

Operational Models of Infrastructure Resilience

David L. Alderson,* Gerald G. Brown, and W. Matthew Carlyle

We propose a definition of infrastructure resilience that is tied to the operation (or function) of an infrastructure as a system of interacting components and that can be objectively evaluated using quantitative models. Specifically, for any particular system, we use quantitative models of system operation to represent the decisions of an infrastructure operator who guides the behavior of the system as a whole, even in the presence of disruptions. Modeling infrastructure operation in this way makes it possible to systematically evaluate the consequences associated with the loss of infrastructure components, and leads to a precise notion of “operational resilience” that facilitates model verification, validation, and reproducible results. Using a simple example of a notional infrastructure, we demonstrate how to use these models for (1) assessing the operational resilience of an infrastructure system, (2) identifying critical vulnerabilities that threaten its continued function, and (3) advising policymakers on investments to improve resilience.

KEY WORDS: Attacker–defender; infrastructure; optimization; resilience; system operation

1. INTRODUCTION

The United States has recently suffered repeated disruptions of its national infrastructure from natural disasters (e.g., Hurricane Katrina in 2005, Superstorm Sandy in 2012), accidental failures (e.g., the Northeast Blackout of 2003), and intentional attack (e.g., World Trade Center and Pentagon attacks of September 11, 2001). In response to these events and to the perceived threat of future ones, the U.S. government has identified 16 critical infrastructure and key resource (CI/KR) sectors.⁽¹⁾ The term “critical infrastructure” is defined in the USA Patriot Act of 2001⁽²⁾ to mean “systems and assets, whether physical or virtual, so vital to the United States that the incapacity or destruction of such systems and assets would have a debilitating impact on security, national economic security, national public health or safety, or any combination of those matters.”

Operations Research Department, Naval Postgraduate School, Monterey, CA, USA.

*Address correspondence to David L. Alderson, Operations Research Department, Naval Postgraduate School, Monterey, CA 93943, USA; dlalders@nps.edu.

Presidential Policy Directive 21 (PPD21) summarizes the government’s objective with regard to critical infrastructure: “The Federal Government also has a responsibility to strengthen the security and resilience of its own critical infrastructure, for the continuity of national essential functions, and to organize itself to partner effectively with and add value to the security and resilience efforts of critical infrastructure owners and operators.”⁽¹⁾ In PPD21, the term “resilience” is defined explicitly to mean “the ability to prepare for and adapt to changing conditions and withstand and recover rapidly from disruptions. Resilience includes the ability to withstand and recover from deliberate attacks, accidents, or naturally occurring threats or incidents.”

In this article, we consider the challenges associated with assessing and improving the operational resilience of critical infrastructure systems. The term “operational resilience” was introduced in an earlier policy document⁽³⁾ in the context of needing to “make the system better able to absorb the impact of an event without losing the capacity to function.” We adopt this term explicitly to mean *the ability of a system to adapt its behavior to maintain continuity*

| Report Documentation Page | | | Form Approved OMB No. 0704-0188 | |
|--|------------------------------------|---|------------------------------------|---------------------------------|
| <p>Public reporting burden for the collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to a penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.</p> | | | | |
| 1. REPORT DATE 2015 | 2. REPORT TYPE | 3. DATES COVERED 00-00-2015 to 00-00-2015 | | |
| 4. TITLE AND SUBTITLE Operational Models of Infrastructure Resilience | | 5a. CONTRACT NUMBER | | |
| | | 5b. GRANT NUMBER | | |
| | | 5c. PROGRAM ELEMENT NUMBER | | |
| 6. AUTHOR(S) | | 5d. PROJECT NUMBER | | |
| | | 5e. TASK NUMBER | | |
| | | 5f. WORK UNIT NUMBER | | |
| 7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Naval Postgraduate School, Department of Operations Research, Monterey, CA, 93943 | | 8. PERFORMING ORGANIZATION REPORT NUMBER | | |
| 9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES) | | 10. SPONSOR/MONITOR'S ACRONYM(S) | | |
| | | 11. SPONSOR/MONITOR'S REPORT NUMBER(S) | | |
| 12. DISTRIBUTION/AVAILABILITY STATEMENT Approved for public release; distribution unlimited | | | | |
| 13. SUPPLEMENTARY NOTES | | | | |
| 14. ABSTRACT We propose a definition of infrastructure resilience that is tied to the operation (or function) of an infrastructure as a system of interacting components and that can be objectively evaluated using quantitative models. Specifically, for any particular system, we use quantitative models of system operation to represent the decisions of an infrastructure operator who guides the behavior of the system as a whole, even in the presence of disruptions. Modeling infrastructure operation in this way makes it possible to systematically evaluate the consequences associated with the loss of infrastructure components, and leads to a precise notion of ???operational resilience??? that facilitates model verification, validation, and reproducible results. Using a simple example of a notional infrastructure, we demonstrate how to use these models for (1) assessing the operational resilience of an infrastructure system, (2) identifying critical vulnerabilities that threaten its continued function, and (3) advising policymakers on investments to improve resilience. | | | | |
| 15. SUBJECT TERMS | | | | |
| 16. SECURITY CLASSIFICATION OF: | | 17. LIMITATION OF ABSTRACT Same as Report (SAR) | 18. NUMBER OF PAGES 25 | 19a. NAME OF RESPONSIBLE PERSON |
| a. REPORT unclassified | b. ABSTRACT unclassified | c. THIS PAGE unclassified | | |

of function (or operations) in the presence of disruptions.

In response to the growing literature on infrastructure resilience, we describe a specific set of analytical tools based on quantitative models of system operation. Specifically, we consider the perspective of the analyst who is charged with (1) assessing the operational resilience of an infrastructure system, (2) identifying critical vulnerabilities that threaten its continued function, and (3) advising policymakers on investments to improve resilience. We present an analysis technique based on the use of a prescriptive model that represents the decisions of an infrastructure operator. That model could be an optimization model, an identity simulation of operating protocols, a heuristic algorithm that mimicks a human operator's decisions, or one of any number of other quantitative tools that can help determine how to operate a system, even in the presence of disruptions. This technique requires that we capture the essential domain-specific details about the infrastructure system in terms of its operator's goals and the limitations on its capabilities. This also requires that we have an unambiguous measure of system performance for the infrastructure. While such features are often not present for general problems in national security and defense, we elaborate on the special features of infrastructure systems that make this technique well suited. To assess the worst-case disruptions to infrastructure function and to identify the most effective defensive measures against them, we apply the game-theoretic *attacker-defender* and *defender-attacker-defender* modeling techniques introduced by Brown *et al.*^(4,5) We illustrate the technique with a simple example and provide mathematical details in the appendices.

A main objective of this article is to advocate in favor of “operational” models that capture domain-specific details relevant to the operation of an infrastructure system. Our intent is not to replace current definitions of resilience; most existing definitions capture some of the essential aspects of resilience, but with very few exceptions they neither provide quantitative (and definitely not operationally based) *measures* of resilience, nor do they provide models that can be used to *improve* resilience. Our primary contribution in this article is to enhance these definitions by making them more precise, and by providing quantitative models that are tied to the performance of the systems in a way that is of direct relevance to the owners and operators of these systems. We hope that these examples—worked out in detail with our

definitions, assumptions, mathematical models, and solution algorithms—will contribute analytical support to the practice of assessing resilience and thus enhancing infrastructure protection.

We develop our models sequentially over the next three sections of the article. In Section 2, we discuss the central Operator Model. Section 3 embellishes the Operator Model to create the Attacker Model, which identifies and evaluates the main vulnerabilities in a system and that can be used both to assess the potential damage to a system due to a set of possible attacks and to define the resilience of a system to a set of attacks. In Section 4, we discuss the Defender Model, which has both the Attacker Model and the Operator Model as subproblems and that can identify optimal, budget-limited ways to improve the resilience of the system to such attacks.

2. IMPORTANCE OF MODELING THE OPERATION OF INFRASTRUCTURE SYSTEMS

Our view of critical infrastructure systems holds that the function of each system, and especially continuity of that function, is of primary importance. In this article, we view an infrastructure as a collection of interconnected *components* that work together as a system to achieve a particular, domain-specific function. It does this through either human or automated decision making that responds to the demands placed on the system to provide the best possible function in any given situation. This decision making is commonly termed the *operation* of the system, and an *operational model* of a system is any mathematical model that evaluates the performance of a system (through a cost function, or some other quantitative evaluation of its operation) and that explicitly includes this operational decision making in its formulation. Although “infrastructure function” in a broad sense may be ambiguous, the notion of function for any particular infrastructure system is typically well defined and understood by its owners, operators, users, and regulators, who develop domain-specific operational models of system performance. For example, the function of an electric power transmission grid (consisting of generators, high-voltage transmission lines, transformers, etc.) is commonly defined by an industry-standard “optimal power flow model” or a related electrical-engineering model (e.g., see p. 419 of Wood and Wollenberg⁽⁶⁾) that determines how well power is being delivered.

Following this, the *importance* of a single component within an infrastructure system is based on how it contributes to the overall function of that system, which we assess as follows. We use the term *disruption* to mean *the loss of one or more system components*, and we measure the *consequence* that results from a disruption in terms of the subsequent loss of system functionality. We calculate this using the operational model to evaluate the change in system performance after the disruption. Having an operational model that provides a clear measure of system function allows us to systematically evaluate the importance of components by considering the consequence associated with their loss, but this requires that we assess how the infrastructure system will respond to each disruption.

In general, the contribution of a single component to system function may depend on its interactions with other components. For example, the loss of a single component might not result in any change to system function (because there is redundancy elsewhere), but the simultaneous loss of this component in combination with other (supposedly) redundant components might be catastrophic to the system. As a result, it is typically not possible to assign a single unique numerical value to each component. Moreover, attempts to rank infrastructure components in terms of such numerical values are certain to be misguided because there might not be a single most-important one (see Alderson *et al.*⁽⁷⁾ for a detailed discussion). Instead, it is more appropriate to discuss the value of *sets of components* that characterizes how important each individual is to the payoff generated by a coalition of players, but applied to system components instead of players, and assessing this is considerably more complicated. In concept, we seek something similar to the “Shapley value” in *n*-person cooperative games⁽⁸⁾ that characterizes who important each individual is to the payoff generated by a coalition of players, but applied to system components instead of players.

We caution against the use of simple surrogate measures of component value (such as replacement cost, or historical importance, or how “connected” a component is to other components), as these measures are far too coarse to indicate a component’s contribution to function, and therefore only have an indirect relationship to system function. Even in simple contexts, such as *maximum flow network problems*,⁽⁹⁾ the most important component (i.e., the component whose loss maximally degrades the flow in the system) is not necessarily the one

with the largest capacity or the one that carries the most flow; in general, these intuitive and appealing approximations do not work.⁽⁷⁾

We also caution against the use of simple surrogate models of system function unless those surrogates are validated against industry-standard models of performance. Over the last decade, there has been a large body of work devoted to the development of purely topological models of infrastructure systems that capture network structure, but little else.⁽¹⁰⁾ For example, some researchers model the function of an electric power grid using graph-theoretic models that emphasize connectivity measures but ignore the physics of electricity transmission, as governed by capacity, inductance, phase angles, etc.^(11,12) Our view is that these topological models fall short of capturing essential domain-specific details needed to represent the operation of an infrastructure system. This view is substantiated by Hines *et al.*,⁽¹³⁾ who show that “evaluating vulnerability in power networks using purely topological metrics can be misleading.” Similar observations have been made for topological models of the Internet.^(14,15)

2.1. An “Operational” View of Infrastructure

The Department of Homeland Security (DHS) states that roughly 85% of the critical infrastructure systems in the United States is owned or operated by the private sector.⁽¹⁶⁾ The behavior of these infrastructure systems is not arbitrary, but reflects an organization that is fundamentally driven by constraints that are placed on their functionality.⁽¹⁷⁾ For example, there are often functional requirements on the system as a whole (e.g., it needs to “work”), which are often stated as objectives (e.g., minimize unmet demand) and then measured in terms of system function. For the private sector, these objectives often take the form of “minimize cost” or “maximize profitability.” In addition, the behavior of the infrastructure is limited by what is possible, due to physical, economic, or regulatory constraints.

In practice, modern infrastructure systems involve a mix of humans (e.g., owners, operators, managers) and autonomous “agents” (e.g., monitoring systems, feedback controllers) that make decisions to guide the behavior of the system as a whole. For example, in California’s electric power infrastructure, the independent system operator (ISO) makes real-time decisions about where to “spin up” or retire generators and which switches to open and close in

the transmission grid so as to route power flow in order to balance generation and demand, subject to constraints on the capacity of individual high-voltage transmission lines and the physics of electricity.⁽¹⁸⁾ The ISO is aided by sophisticated supervisory control and data acquisition (SCADA) systems that implement decision rules for managing the system as operating conditions change.

We refer to this collective decision-making entity as “the operator” of the infrastructure. Some infrastructure systems have explicit operators (e.g., electric power), while others are governed by the interaction of many decision agents (e.g., drivers of vehicles in a regional road system). In the latter case, we can often represent the collective decision-making behavior in terms of an equilibrium model.^(19,20)

The key point is that the operator makes *decisions* about the behavior of the system in order to reconcile these *objectives* (what we want the system to do) with its *constraints* (what the system can do) in an intelligent manner. The language of *constrained optimization* is ideally suited to represent this type of decision problem (see Rardin⁽²¹⁾ for an introduction), and we adopt constrained optimization hereafter, though other types of models such as simulations might apply in other contexts.

Optimization models of this type are *prescriptive*: potential courses of action are represented using *decision variables*, and the solution to a particular problem indicates decisions that *should be taken* to reconcile objectives and constraints in a best possible manner (where “best” reflects the stated objective).

Modeling the behavior of an infrastructure system in terms of a constrained optimization problem does not necessarily mean that we believe that the real operation of the system is truly optimal. Rather, *the key to a “good” operational model of infrastructure is to identify the essential structural features, defined in terms of the problem’s objectives and constraints*. We make several arguments in support of this claim. First, the solution to a constrained optimization problem that more less gets an infrastructure’s basic objectives and constraints correct is going to display behavior that looks a lot more realistic than a model of behavior that completely ignores system function, operating objectives, and constraints.^(10,17,22) This will be the case even if the model solutions are only near-optimal, and even an approximate solution to a constrained optimization problem can provide insight into infrastructure behavior. Second, real infrastructure owners and operators regularly formulate and solve constrained

optimization problems to guide their decisions about how to run their systems. Often, there are industry-standard models of infrastructure that can be adopted as realistic representations of infrastructure behavior. We advocate using such models whenever available. Third, there is now a large literature in operations research devoted to formulating and solving these problems. Recent advances in mathematics and computation allow us to solve problems of realistic scale and fidelity in this manner.

There is one other key advantage to using an optimization-based prescriptive model of system operation as the starting point for the study of infrastructure behavior: *these models naturally accommodate disruptions to infrastructure as straightforward changes to input data*. For example, Salmerón *et al.*⁽²³⁾ present a model of electric power transmission that takes available generators, transmission lines, transformers, and buses, and identifies the set of power flows that minimizes unmet, prioritized demand; this model has been validated as a realistic representation of the actual grid. If the system loses a transformer, we would like to know: How will this system adapt its behavior, and what will be the consequences on system function? We simply need to re-solve the same operator’s problem, leaving the affected transformer “out” of the model (how exactly this happens will depend on the implementation of the model, but it is essentially an input modification); then, the solution to this modified problem will indicate the best possible response of the system. Thus, *system adaptation is inherent to the model formulation, not an afterthought*.

This basic form of a decision model for the operator offers exactly what we need to systematically evaluate the consequences associated with the loss of sets of components. For example, we can investigate specific disruption scenarios of interest by rerunning the same model to find the best response to each. However, because we have defined the set of possible disruptions in terms of the loss of components, we can also consider a broader evaluation of all possible disruptions (e.g., via exhaustive enumeration).

Observe that you could never do this with a model that is purely *descriptive* (e.g., via a set of differential equations that describe *a priori* all future states of the system) because it would require that you consider in advance all of the possible contingencies in disruption and response, and account for them in the predefined description of behavior. Thus, the use of a prescriptive model has the benefit of not

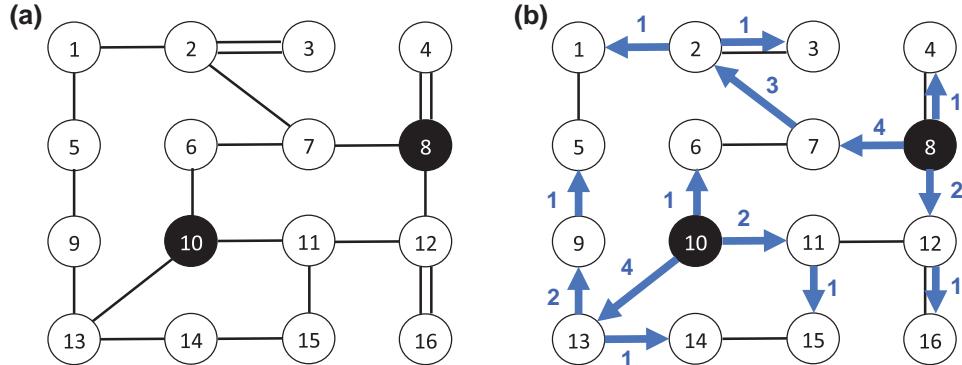


Fig. 1. A notional infrastructure system. (a) A white circle (node) represents a location with demand equal to one barrel of fuel. A black circle (node) represents a location with supply equal to 10 barrels. Each link is bidirectional, has a fuel flow capacity of 15 barrels, and has per-barrel transit cost of \$1. The penalty for unsatisfied demand per node is \$10 per barrel. Nodes 3, 4, and 16 each have two (parallel, redundant) connections to the rest of the network. This network has been built to be $N-1$ reliable, meaning that the loss of any single link does not disconnect any node. (b) Shows baseline flows corresponding to a minimum-cost flow solution, which results in a total cost of \$25.

needing to specify the entire “trajectory” for system behavior (or set of possible trajectories) in advance. Rather, when there is a disruption, one simply solves for the best course of action going forward. This is more in line with what real infrastructure owners and operators do in practice.

2.2. A Notional Example

Consider a notional infrastructure system designed to distribute some commodity, say, *fuel*, to different locations within a city (the metaphor here is a simplified petroleum distribution system, but the modeling technique is general). Fig. 1(a) presents a simple distribution network between two supply locations (represented by black *nodes*) and 14 demand locations (represented by white nodes). Fuel is carried by *links* that are bidirectional (meaning that flow can move in either direction) and have a limited flow capacity. Assume that the demand for fuel at each demand location is one *barrel* of fuel, that the supply of fuel at each storage location is 10 barrels, and that each link can carry up to 15 barrels of fuel.

The *operator* of this infrastructure system makes decisions about how to manage fuel flows based on costs. Specifically, assume the operator faces a contractual penalty of \$10 per barrel for each location that does not receive its demanded fuel. In addition, assume that the per-unit cost to send fuel over a single link is \$1 per barrel.

The operator’s objective is to route the available fuel so as to minimize the sum of all delivery costs

and penalty costs for the system. This task is complicated by the fact that one or more of the links in this system can be *broken* (equivalently, *failed*, *lost*, *attacked*, or *interdicted*). The operator faces the same objective even when there are broken links in the system—in this case, she must do the best she can to minimize the sum of delivery costs and penalties with the surviving distribution network.

We define the *Operator Model* as a constrained optimization problem of the following form:

$$\min_{\mathbf{y} \in Y(\hat{\mathbf{x}})} f(\hat{\mathbf{x}}, \mathbf{y}), \quad (1)$$

where $\hat{\mathbf{x}}$ is a vector that collectively represents whether each of the components (the links in our example) in the system is working or broken (also called the operating *state*), the set $Y(\hat{\mathbf{x}})$ represents the feasible actions of the operator (here, allowable flows) for given state $\hat{\mathbf{x}}$ of the system, and $f(\hat{\mathbf{x}}, \mathbf{y})$ is a function that measures the performance (here, the cost) that results from the choice of activities \mathbf{y} . The operations research literature is filled with such models, although most do not explicitly parameterize damage. Appendix A presents a formal mathematical representation of the Operator Model for this example.

Given the potential for broken links, the network in Fig. 1(a) has been constructed so as to be $N-1$ reliable (a standard notion in system reliability, where N denotes the total number of system components), meaning that a single broken link cannot disconnect any node in the network. In particular, there are two sources of fuel, and three of the locations (labeled

as nodes 3, 4, and 16) are each connected by parallel links so a single break does not disconnect them.

Figure 1(b) shows the minimum-cost flows to deliver fuel to each location when there is no broken link; this is the baseline solution to the Operator Model. This system is balanced and has excess capacity—each of the sources are supplying 50% of the total demand (7 of 14 units demanded), and each has 30% reserve storage beyond what is delivered (using 7 of 10 units of available fuel).

3. ASSESSING THE RISK OF POSSIBLE DISRUPTIONS

An operational model of infrastructure behavior allows us to systematically evaluate how the system will respond to any disruption (defined in terms of the simultaneous loss of one or more system components) and then measure the consequence in terms of a change in system function. The key question becomes: What kinds of disruption scenarios are of most concern?

3.1. Nondeliberate Hazards Versus Deliberate Threats

In practice, infrastructure owners and operators must contend with both *nondeliberate hazards* (e.g., accidents, failures, and Mother Nature) and *deliberate threats* (e.g., vandalism, sabotage, competitors, and terrorism). The study of failures in technological systems has yielded an extensive literature on *system reliability*.^(24,25) The broader study of risk in the context of nondeliberate hazards has resulted in a large literature in *probabilistic risk analysis (PRA)* that defines possible future scenarios, assigns a probability to each scenario, estimates the consequence associated with each scenario, and then aggregates this information into one or more measures of risk, such as expected value, value at risk,⁽²⁶⁾ or conditional value at risk.⁽²⁷⁾ PRA has been particularly successful when applied to nondeliberate hazards for which there are data or models that can be used to assess the required probabilities. In some cases, these data may be historical (e.g., weather records, failure statistics, actuarial statistics, and accident reports) or can be obtained via experiment (e.g., laboratory stress testing to evaluate the mean time between failures). For so-called rare events there is ongoing debate about how to model the frequencies with which disruptions occur (e.g., earthquakes^(28,29)), and this is an active area of research.

Following the attacks of September 11, 2001, there was a shift in national priority from assessing nondeliberate hazards to preventing and protecting against deliberate threats, and the study of risk in national security problems has been controversial ever since. Paté-Cornell and Guikema⁽³⁰⁾ are among the first to apply the techniques of PRA to terrorism risk. Many papers follow,^(31–35) often using simplified models that rely on the definition “Risk = Threat (T) \times Vulnerability (V) \times Consequence (C),” where subject matter experts assess the threat and vulnerability terms as probabilities, and the consequence term in units of, for example, economic replacement cost, or fatalities.^(36,37) When applied to critical infrastructure, the notion is to assess adversary intent as “threat”⁽³⁸⁾ and then rely on such assessments for proposed methods to optimize defense.^(39,40) DHS has promoted PRA, including models based on the (T,V,C) construct, for assessing the threats posed by intelligent adversaries in a terrorist attack.⁽⁴¹⁾

The National Research Council (NRC) has, however, criticized the use of probabilities to model the behavior of an intelligent, goal-oriented terrorist.^(42,43) Additional work has raised concerns about terrorism risk models based on (T,V,C). For instance, with a number of examples, Cox⁽⁴⁴⁾ illustrates how these models can render nonsensical advice. Cox⁽⁴⁵⁾ further notes the deficiency of T, V, and C values as inputs when the probabilities are correlated, and Cox⁽⁴⁶⁾ also points out that because the values for V and C really depend on the allocation of effort by both the attacker and defender, they do not make sense as independent inputs. Brown and Cox^(47,48) detail several ways in which probabilistic assessment of terrorism risk can mislead analysts, and they explain why it is impossible for a defender to possess information essential to assess terrorist intent.

Without revisiting the arguments on both sides of this debate, we comment on a few issues most relevant to the resilience of infrastructure systems. *First and foremost, if using a (T,V,C)-style of analysis for an infrastructure system, one cannot assume that the consequence associated with the loss of a set of components is simply the sum of the consequences associated with the loss of individual components.* In general, ignoring the dependencies between the components of a system can be misleading. Rather, one should be considering scenarios involving the loss of sets of components. In some risk analyses, the “components” of the system are themselves built of

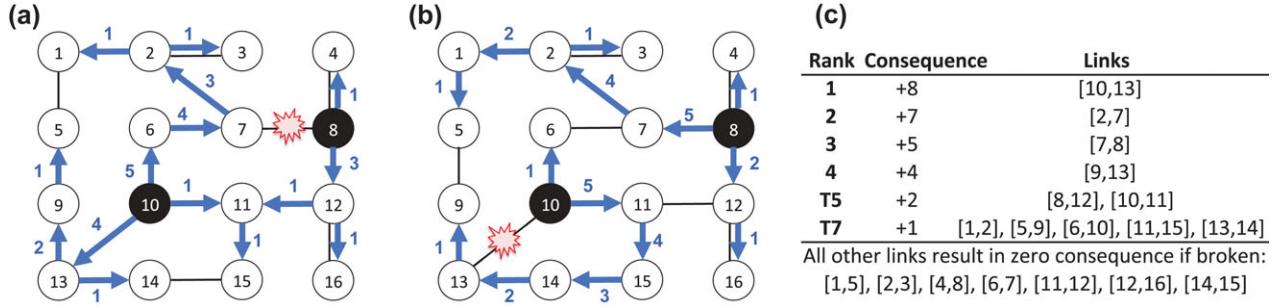


Fig. 2. A break on a single link in this network incurs additional operating cost, but does not prevent fuel from being delivered to each location. (a) A break on link [7, 8] results in an increased operating cost of 32. (b) A break on link [10, 13] is the worst-possible interdiction of a single link and results in total cost of 33; in this case, there are multiple ways the operator can reroute flows and achieve this cost. (c) This table lists the links, if interdicted individually, that yield the greatest consequence, in rank order. “T5” and “T7” denote ties for fifth-worst and seventh-worst, respectively.

elements, modeled, and evaluated by any of a number of probabilistic models.^(49,50) But this is not current practice in many implementations of PRA for critical infrastructure systems. We therefore caution against the use of simplistic (T,V,C)-style modeling for the study of deliberate threats to critical infrastructure, and especially for assessing infrastructure resilience.

Second, our operational view of infrastructure function is agnostic to the source of a disruption—once one or more components are lost, the operator’s focus is on doing the best she can to maintain function with whatever is left of the system.

Third, by narrowing the set of potential disruptions to the simultaneous loss of one or more (known, and finite) components, it becomes possible, in principle, to search over the scenarios of concern. Although the size of this set can be too large to allow this in practice, it creates the opportunity for a different style of analysis, as we now describe.

3.2. Using the Operator Model to Assess Disruptions

Given the Operator Model (Equation (1)), we can explicitly consider the consequence of any potential disruption (i.e., loss of links in our notional infrastructure) by changing \hat{x} and re-solving for the minimum-cost response. For example, consider a break in the link [7, 8] as shown in Fig. 2(a). This link previously carried 40% of the total system flow in the baseline solution. In response to this break, and under the assumptions of this example, the system operator is able to reroute flows through the network in

order to still satisfy all customers; however, the total cost to do so increases from 25 to 32.

Although the network is $N - 1$ reliable, suppose the operator is concerned about the worst-case loss (break) of a single link because it will create the need to reroute flows and possibly incur greater cost. One way to find the worst single-link loss in the system is to exhaustively enumerate each possible interdiction, re-solving Equation (1) each time, and then identifying the possible interdiction that results in the highest operating cost.

Another way to get at this is to consider a hypothetical intelligent adversary (an *attacker*) who has perfect knowledge of the system and uses limited resources to deliberately damage the system. From the operator’s perspective, the attacker could be Mother Nature, a terrorist, simple bad luck, or anything else that causes the simultaneous loss of components; the operator is concerned with running the system in the best possible manner following the loss of these components. Although our exposition sometimes personifies the attacker, we emphasize that our purpose is simply to discover worst-case component losses, not model the actual decision making of any particular adversary (e.g., Al-Qaeda).

Suppose the attacker has the ability to target a single link. Which one should he break to maximize the costs incurred by the operator? We formulate this *Attacker Model* mathematically as follows:

$$\max_{\mathbf{x} \in X} \min_{\mathbf{y} \in Y(\mathbf{x})} f(\mathbf{x}, \mathbf{y}), \quad (2)$$

where now \mathbf{x} is a decision variable belonging to the attacker, and X represents the set of all possible single-link attacks. Given any particular choice

of attack \mathbf{x} , the operator still faces the same cost minimization (Equation (1)), now with an objective function $f(\mathbf{x}, \mathbf{y})$ and a set of feasible actions $\mathbf{y} \in Y(\mathbf{x})$. Thus, this Attacker Model (Equation (2)) is almost identical to the prior model (1), except that the state parameters \hat{x} have become decision variables for the attacker, and we have put restrictions on the choice of disruption. Models of this form have been studied in the context of *attacker–defender optimization*.^(5,7,51)

Appendix B provides a complete mathematical formulation for the Attacker Model in this example. For our notional infrastructure, the worst-case single-link disruption is the loss of link [10, 13], which results in a total operating cost of 33 (Fig. 2(b)). The table in Fig. 2(c) lists the links, if interdicted individually, that yield the greatest consequence, in rank order.

An important contribution in the development of attacker–defender optimization problems is their connection to game theory.⁽⁵²⁾ Specifically, the mathematical program (Equation (2)) is a two-stage sequential-play game in which the attacker moves first, and then the operator (or defender) moves second. These are known as *Stackelberg games*.⁽⁵³⁾

If all of the decision variables for the attacker and defender are discrete, our formulation (2) is equivalent to a sequential matrix game of the classical layout, where in the first stage the attacker chooses a row of the payoff matrix by choosing a particular attack plan, and then the operator (or defender) chooses a column through his choice of a specific operating plan. However, instead of enumerating all of the pure strategies for each player at each stage of the game, we represent those (potentially enormous) sets of pure strategies implicitly through a set of decision variables and constraints: the exponential number of feasible solutions to this constrained optimization model represent the possible pairs of strategies for the two players.

This implicit representation of the strategy spaces allows us a great deal of power in modeling the behavior of the two players. We can impose any number of budget restrictions on each player (e.g., time, money, labor, explosives, or other materials), and we can also add constraints that preclude illogical (or physically impossible) combinations of decisions, and in this way we can represent extremely complex decision spaces with only slightly more modeling effort.

The ability to solve attacker–defender problems in this manner also has implications on how we assess the resilience of such infrastructure systems.

4. ASSESSING AND IMPROVING RESILIENCE

Resilience has recently become an important topic in discussions about the way that systems of all kinds respond to both nondeliberate hazards and deliberate threats. This section describes how our operational view of infrastructure function naturally leads to a precise and quantifiable notion of “operational resilience,” and we describe how our optimization-based attacker–defender models lend themselves to characterizing it in a way that facilitates model verification, validation, and reproducible results—features that are essential to making the study of resilience more scientifically rigorous.

4.1. Notions of Resilience

Park *et al.*⁽⁵⁴⁾ provide a partial survey and summary of the growing literature on resilience and its relationship to the study of risk. They report how use of the term “resilience” in engineering systems followed the foundational work of Holling⁽⁵⁵⁾ in ecology, with considerable growth in the number of papers in the last decade that relate to resilience in engineering, physics, and mathematics. Hollnagel *et al.*⁽⁵⁶⁾ provide an early treatment of “resilience engineering” that builds on the study of system safety. As noted by Madni and Jackson,⁽⁵⁷⁾ an important contribution in this early work is the argument that “safety is something (that results from what) a system or an organization *does*, rather than something a system or an organization *has*.” As a result, much of this literature stresses the need to study safety as a process instead of safety as a property of the system itself. The study of resilience in engineering systems has followed this lead, in the sense that resilience is viewed as an expression of system behavior in response to an event rather than something inherent to the system itself.^(57,58)

A complicating factor in previous attempts to define resilience is the recognition that “[r]esilience is a family of related ideas, not a single thing.”⁽⁵⁹⁾ Zolli and Healy⁽⁶⁰⁾ provide perhaps the most comprehensive and provocative discussion of the myriad notions of resilience. Nonetheless, a common feature across many definitions of engineering resilience is the ability of the system to *adapt* in response to a disruption.^(57,58,61–63) Importantly, Park *et al.*⁽⁵⁴⁾ observe that in a resilient system the result of this adaptation is “the persistence of relationships, rather than stability in quantitative measures of state variables.”

Thus, a distinguishing feature of resilience is adaptation in the way that components work together to achieve persistence in these relationships. Our notion of operational resilience is consistent with these ideas in the sense that our focus is persistence in the ability of a system to function, over time, in the presence of disruptions.

Park *et al.*⁽⁵⁴⁾ further comment on why resilience in engineering systems should be different from that in ecology, and why it is distinct and complementary from the study of risk. They argue that the emergent, nonlinear, self-organizing features in coupled complex systems make hazard identification difficult if not impossible, that assessing the probabilities of harm may be unknowable, and that “we have a poor understanding of how failures propagate and amplify within and across complex systems.” Although we agree with the notion of resilience “not as something a system has, but a characteristic of the way it behaves,” we take issue with the claim that engineering resilience in a system “cannot be predicted or calculated from aggregation of the individual system components.”⁽⁵⁴⁾ Modern infrastructure systems are complicated, and they can also exhibit features of complexity (see Ottino⁽⁶⁴⁾ for a discussion of the distinction between “complicated” and “complex”); however, designating an infrastructure as a “complex system” does not mean that we are at the mercy of nonlinear, emergent chaos or self-organization. Rather, the fundamental belief underlying our Operator Model is that by capturing the essential objectives and constraints driving system behavior, we build a representation that is explanatory and not merely descriptive, in the sense of Willinger *et al.*,⁽⁶⁵⁾ and that this representation will therefore have superior predictive power for assessing the “what-ifs” associated with disruption.

In the last decade, there have been considerable efforts within the engineering community to assess the resilience of infrastructure systems. Haimes *et al.*⁽⁶⁶⁾ observe that “[o]ne approach to measuring the resilience of an infrastructure is to predict the trajectory of recovery time following a catastrophic event.” Reed *et al.*⁽⁶⁷⁾ present resilience scoring metrics and build on the work of Haimes⁽⁵⁸⁾ in using input-output models to measure the resilience of interconnected systems. These ideas have been prevalent in the civil engineering literature, particularly in assessing the resilience of freight transportation^(68,69) and its dependence on maritime systems,^(70,71) with emphasis to evaluate the resilience of transportation networks after a disaster.^(72,73)

Using ideas from control theory, Vugrin *et al.*⁽⁷⁴⁾ characterize resilience in terms of the deviation (both magnitude and duration) from “normal” operation that follows a disruptive event; in this context, a system is more resilient if it experiences smaller deviations. Vugrin *et al.*⁽⁷⁵⁾ use their definition to assess the resilience of the U.S. petrochemical sector in response to two hypothetical hurricane scenarios in the Gulf Coast region. Rose⁽⁷⁶⁻⁷⁸⁾ has studied economic resilience to disasters in terms of distinct phases of service restoration and economic recovery over time.

Despite recent efforts to develop common resilience metrics across infrastructure systems, Haimes⁽⁵⁸⁾ cautions against the use of scoring for system resilience: “attempts to characterize the resilience of a system with a specific numerical descriptor (as a metric) and to use the metric to compare the resilience of different systems could be misleading” because of the differences in operating environments for different infrastructure systems.

As noted, resilience has become an important concept in discussions about homeland security and defense. A March 2010 report by the U.S. Government Accountability Office (GAO)⁽⁷⁹⁾ traces the history in the definition and use of resilience in the U.S. government’s official documents on homeland security and also details the increased role of resilience in the updated 2009 National Infrastructure Protection Plan. DHS currently defines resilience as the “ability to adapt to changing conditions and prepare for, withstand, and rapidly recover from disruption.”⁽⁸⁰⁾ Outside the government, Flynn⁽⁸¹⁾ points to vulnerabilities that threaten our national welfare and provocatively asks how the United States can rebuild itself into a more resilient nation.

Despite this recent flurry of activity, a key challenge remains how to define resilience in a manner that is (1) quantitative and rigorous enough for objective and precise assessment, (2) flexible enough to capture many facets of resilience already under discussion by researchers, and (3) connected to the operational details of the system under study so that proposed system changes can be naturally evaluated and actually implemented. We proceed in direct support of this objective.

4.2. Assessing Operational Resilience

Resilience is fundamentally about the behavior of a system in response to a disruption. Our focus on infrastructure systems and use of an Operator

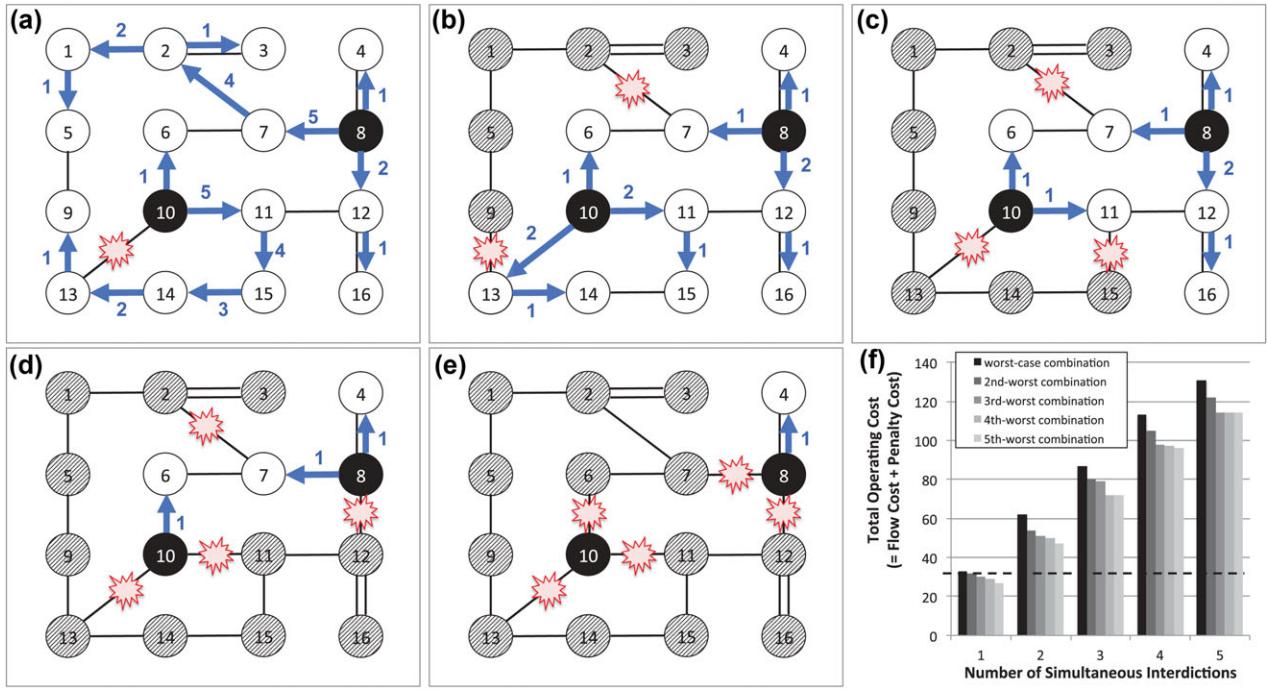


Fig. 3. Worst-case simultaneous interdictions. (a) The worst-case single interdiction is of link [10, 13], resulting in a total cost of 33. In this case, the flow cost increases but all nodes are still served. (b) The worst-case simultaneous two-link interdiction is of links [2, 7] and [9, 13], which denies nodes 1, 2, 3, 5, and 9 (now shaded) any flow. The total cost is 62 ($=12 + 50$), most of which is unmet demand penalty cost. (c) The worst-case simultaneous three-link interdiction is of links [2, 7], [10, 13], and [11, 15], resulting in a total cost of 87 ($=7 + 80$). (d) The worst-case simultaneous four-link interdiction is of links [2, 7], [8, 12], [10, 11], and [10, 13], resulting in a total cost of 113 ($=3 + 110$). (e) The worst-case simultaneous five-link interdiction is of links [6, 10], [7, 8], [8, 12], [10, 11], and [10, 13], resulting in a total cost of 131 ($=1 + 130$). (f) The worst-case (rank 1) attack for 1–5 simultaneous interdictions increases approximately linearly. The second-worst (rank 2) through fifth-worst (rank 5) attacks do less damage, but all are significantly worse than the baseline (no interdiction) case that has operating cost 25.

Model to represent infrastructure behavior requires us to define the system in question, specify its components, and provide an unambiguous measure of system performance. In this section, we show how our definition of *operational resilience*—that is, the ability of a system to adapt its behavior to maintain continuity of function (or operations) in the presence of disruptions—can be assessed in a straightforward manner by performing parametric analysis using our Attacker Model.

Alderson *et al.*⁽⁷⁾ introduce the notion of a *resilience curve* as that which plots the best achievable worst-case performance of a system as a function of the disruption “magnitude” that we measure, for example, in terms of the number of simultaneously lost components. The usefulness of a resilience curve is based on two underlying ideas. First, by classifying disruptions in terms of the number of lost components, we obtain a natural mechanism for considering disruptions that range from “small” to “large.” This

is important in comparing different systems because the way that each responds to disruptions of different sizes can be dramatically different and even make it difficult to say which one is “more resilient” (for a detailed discussion, see Alderson *et al.*⁽⁷⁾). Second, for any particular magnitude of disruption, we conservatively focus on the worst-case loss of components. Thus, our notion of the “worst-case” component loss is always implicitly conditioned on some admissible set of combinations of lost components. Most simply, we often consider the set defined by the maximum number of lost components (i.e., a cardinality constraint), but this generalizes to any notion of “budget” including an explicit list of attack options that are affordable to a specific attacker. We find this parameterization to be of more practical value than the “absolute worst-case,” which reasonably might correspond to the simultaneous loss of *all components*.

With this in mind, consider the worst-case disruption in our notional example associated with the

simultaneous loss of from one to five links (Fig. 3). Specifically, in the presence of the worst-case loss of a single link (Fig. 3(a)), our network is able to reroute flows in order to satisfy demand at all nodes. However, the worst-case loss of two and more links (Figs. 3(b)–(e)) effectively isolates nodes and incurs escalating operating costs (Fig. 3(f)), due primarily to the model penalties for unmet demand. The frontier associated with the worst-case losses of one-to-five links (black bars in Fig. 3(f)) is our “resilience curve” for this example; here, it shows that an attacker can get approximately linear returns for each additional attack.

The relative shape of this “curve” reveals a lot about the resilience of the system. We would say that a system for which operating costs grow more quickly with the number of lost components is “less resilient” than our example, and that a system whose operating costs grow less quickly with the number of lost components is “more resilient.”

We obtain the results in Fig. 3 by solving the Attacker Model (Equation (2)) with a simple constraint on the feasible number of attacks, which we vary parametrically from $k = 1$ to $k = 5$ total attacks (see Appendix B for details). Fig. 3(f) also shows the operating costs associated with the second-worst (i.e., rank order 2) through fifth-worst (rank order 5) combination of losses for each magnitude of disruption. In principle, obtaining these rank-ordered disruptions is no more complicated than exhaustively enumerating each possible loss of k components and then sorting by consequence. However, due to the large number of combinations, in practice it is more efficient to solve the Attacker Model repeatedly, each time with an additional constraint that eliminates the previous solution from further consideration (see example S1 on p. 156 of Brown and Dell⁽⁸²⁾). Discovering the worst, second-worst, third-worst, etc., disruptions has important practical considerations for assessing system resilience and advising defensive investment. If there is only a single unique worst-case disruption with consequence that is much larger than the second-worst, then defending against that single disruption might be sufficient to dramatically increase the resilience of the system. In contrast, if the worst-case disruption is not unique but is accompanied by many equally bad ones, then defending against only one of them is unlikely to help at all.

Thus, an analysis of infrastructure function using the attacker–defender technique leads to a natural characterization of operational resilience.

4.3. Improving Operational Resilience

Our ultimate goal is not just defining and assessing, but improving operational resilience of our infrastructure systems. In the context of our Operator Model, this means mitigating the worst-case operating cost that can result from the simultaneous loss of components. However, doing so will require investment, and our ability to spend on improvements will be constrained by limited resources. To quantify this decision, we formulate this *Defender Model* mathematically as follows:

$$\min_{\mathbf{w} \in W} \max_{\mathbf{x} \in X} \min_{\mathbf{y} \in Y(\mathbf{w}, \mathbf{x})} f(\mathbf{w}, \mathbf{x}, \mathbf{y}), \quad (3)$$

where \mathbf{w} is a decision variable representing defensive investments, and W represents the set of feasible investments. These investments potentially change the operating cost $f(\mathbf{w}, \mathbf{x}, \mathbf{y})$ faced by the operator, as well as the set of feasible actions $\mathbf{y} \in Y(\mathbf{w}, \mathbf{x})$. Models of this form have been studied in the context of *defender–attacker–defender optimization*.^(5,83) Appendix C provides a complete mathematical formulation for the Defender Model in this example.

4.3.1. Protection

We consider two defensive strategies for improving operational resilience. First, assume we have the ability to protect (equivalently, “harden”) a link so that it is invulnerable to loss. For our notional attacker, this means that an attack on the protected component will not affect system performance. In order to identify the worst-case disruption in the presence of protection, we further assume that this attacker can see which links have been protected before he decides what to attack. Given some limited ability to defend links in this way, which links should we protect, and how will this change the worst-case attack and the resulting consequence?

Fig. 4 displays the optimal defenses against a given number of attacks. Each row corresponds to a single link in our notional infrastructure. Each column corresponds to a scenario involving a specified number of defenses and attacks. The column values for each scenario represent the optimal defenses (denoted as “O”) against that number of attacks, as well as the worst-case attacks (denoted as “X”) in response to those defenses. Fig. 5 illustrates in more detail the optimal defenses against the worst-case attack on three links. Here we obtain insight into the strategy for defensive protection—the optimal defense is one that “breaks up the set of

| attacks defenses | 0 | 1 | 2 | 3 | 4 | 5 | 0 | 1 | 2 | 3 | 4 | 5 | 0 | 1 | 2 | 3 | 4 | 5 | 0 | 1 | 2 | 3 | 4 | 5 | | | | | | |
|---------------------|----|----|----|----|----|----|----|----|----|----|----|----|----|----|----|----|----|----|-----|----|----|----|----|----|-----|-----|----|----|----|----|
| edges | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
| [1,2] | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
| [1,5] | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
| [2,3] | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
| [2,7] | X | O | O | O | O | O | X | O | O | O | O | O | X | X | O | O | O | O | X | X | O | O | O | O | | | | | | |
| [4,8] | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
| [5,9] | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
| [6,7] | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
| [6,10] | X | O | O | O | O | O | X | O | O | O | O | O | X | X | O | O | O | O | X | X | X | O | O | O | | | | | | |
| [7,8] | X | O | X | X | X | O | X | X | X | X | X | X | X | X | X | X | X | X | X | X | O | X | X | X | | | | | | |
| [8,12] | X | O | X | O | O | O | X | X | X | X | X | X | X | X | X | X | X | X | X | X | O | X | X | X | | | | | | |
| [9,13] | X | O | O | X | O | O | X | X | X | O | X | O | X | X | O | X | X | O | X | X | O | X | X | O | | | | | | |
| [10,11] | X | | | | | | X | X | O | O | O | O | X | X | X | X | X | X | X | O | O | O | O | O | | | | | | |
| [10,13] | X | O | O | O | O | O | X | X | O | X | X | O | O | O | O | O | O | O | O | O | O | O | O | O | | | | | | |
| [11,12] | | | | | | | X | | | | | | | | | | | | | | | | | | | | | | | |
| [11,15] | | | | | | | X | X | O | X | X | X | X | X | X | X | X | X | X | X | X | X | X | X | | | | | | |
| [12,16] | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
| [13,14] | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
| [14,15] | | | | | | | | | | | | | | | | | | | | | | | | X | | | | | | |
| cost | 33 | 32 | 30 | 29 | 27 | 27 | 62 | 50 | 47 | 47 | 42 | 40 | 87 | 80 | 67 | 64 | 49 | 43 | 113 | 82 | 74 | 71 | 63 | 55 | 131 | 104 | 96 | 80 | 65 | 59 |

Fig. 4. Optimal defensive “hardening” of links can mitigate the worst-case attack. Here, an “O” represents the protection of a link, and an “X” represents an attack. For a given number of attacks, an optimal defense “breaks up” the worst-case set of attacks, and the attacker finds the next-worst set of attacks. The case of three attacks is additionally illustrated in Fig. 5. Scanning across rows here reveals that the links in this network cannot be ranked in a simple priority list of importance; however, the frequency with which a link appears in attack or defense solutions provides an indication of relative importance. The bottom row shows the optimal, postattack operating cost for each scenario.

attacks” that yield the worst-case operating cost. For our notional example, the high penalty costs associated with unmet demand means that the worst-case attack is the one that disconnects as many nodes as possible. The optimal defense prevents this by ensuring that as many nodes stay connected as possible, even in the presence of three interdicted links.

This general pattern is observed throughout Fig. 4. Moving from left right for a fixed number of attacks, each successive column could be interpreted as a type of iterative fictitious play—in which the attacker selects an attack set, the defender protects a link to “break up” the attack set, then the attacker selects a new attack set, after which the defender protects another link to counter that attack set, and so on—that is used to obtain the final solution to that specific combination of attacks and defenses. We emphasize, however, that the actual “game” being played here has only three stages: the defender moves first by protecting some links, the attacker selects the vulnerable links to interdict, and the operator runs the residual system as best she can to minimize operating costs of the surviving system. The solution for each column is obtained by solving an instance of our Defender Model (Equation (3)) with the corresponding number of defenses and attacks.

Fig. 4 also reports the resulting postattack operating cost for each scenario. We illustrate these in Fig. 6 as the resilience curves associated with increased defenses. In the absence of protection, after the first attack, the postattack operating costs (for our simple example) grow approximately linearly with the number of attacks.

In the case where all attacks are equally costly, the resilience curve for an infrastructure can also be viewed as a simplified form of *the attacker’s return on investment (ROI)*. For our simple example, the linear shape of this curve is not good news for the operator. Fortunately for this system, with each additional defense this curve becomes less steep, reflecting the fact that attacks become less effective. Thus, protecting links in this manner improves the resilience of the system—the system denies consequences to the attacker, no matter his actions.

4.3.2. New Construction

Another strategy for creating resilience in an infrastructure system is to augment it with new construction. Specifically for our notional example, assume we have the ability to build any of the dashed-line links shown in Fig. 7(a) and that any

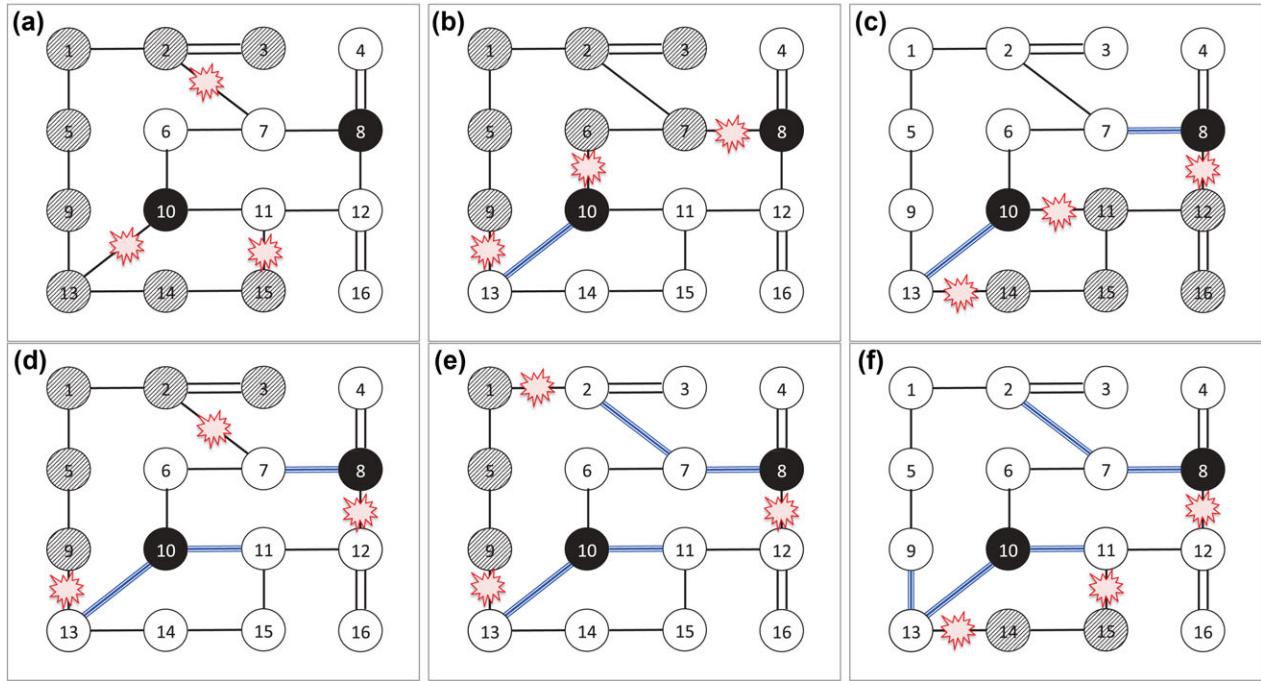


Fig. 5. Protecting links mitigates a worst-case three-link attack. Panels (a)–(f) display a worst-case attack on three links in the presence of 0–5 defenses, respectively. In the case of three attacks, the defensive importance of individual links follows a simple priority list: [10, 13, 7, 8, 10, 11, 2, 7, 9, 13]. With each additional defense, the worst-case attack results in a lower consequence (however, note the associated attack changes completely). The corresponding costs appear in Fig. 4.

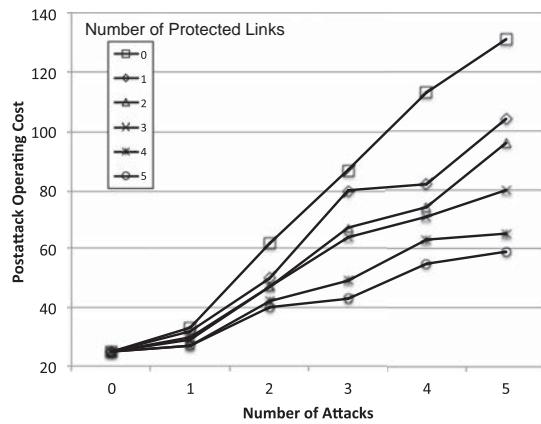


Fig. 6. Resilience curves. In the absence of protection, the postattack operating costs grow approximately linearly with the number of attacks. With each additional protection, this curve (for our simple example) becomes less steep, indicating improved operational resilience for the system.

newly built links will be invulnerable to attack. However, we also assume that it “costs” twice as much to build a new link than to protect an existing one. Under these assumptions, which links, if any, should we build, and which links should we protect?

The table in Fig. 7 shows the postattack operating cost for different defensive budget levels. Here, we represent defensive budget in simple cardinality terms (to simplify exposition—we have included much more complicated investment considerations in other such models), where it “costs” one unit of defense to protect a single link and two units of defense to build a new, invulnerable link. For each budget level, we consider all possible combinations of links to build and protect, and for each combination we solve the Defender Model (Equation (3)) with each assumed number of attacks. The values in this table report the resulting postattack operating costs, and a smaller cost value indicates a better defense.

In many of these cases, we observe that it is more effective to build new links than to defend existing ones. This is not surprising because adding links to the network serves to shorten the average path length between nodes in the network, and this helps to reduce the operating cost of the system, in addition to providing redundant paths. However, building new links is not a strictly dominant strategy, and even in this small example we observe all combinations of build-only, build-some-protect-others, and protect-only. Figs. 7(b)–(e) illustrate in more detail the best

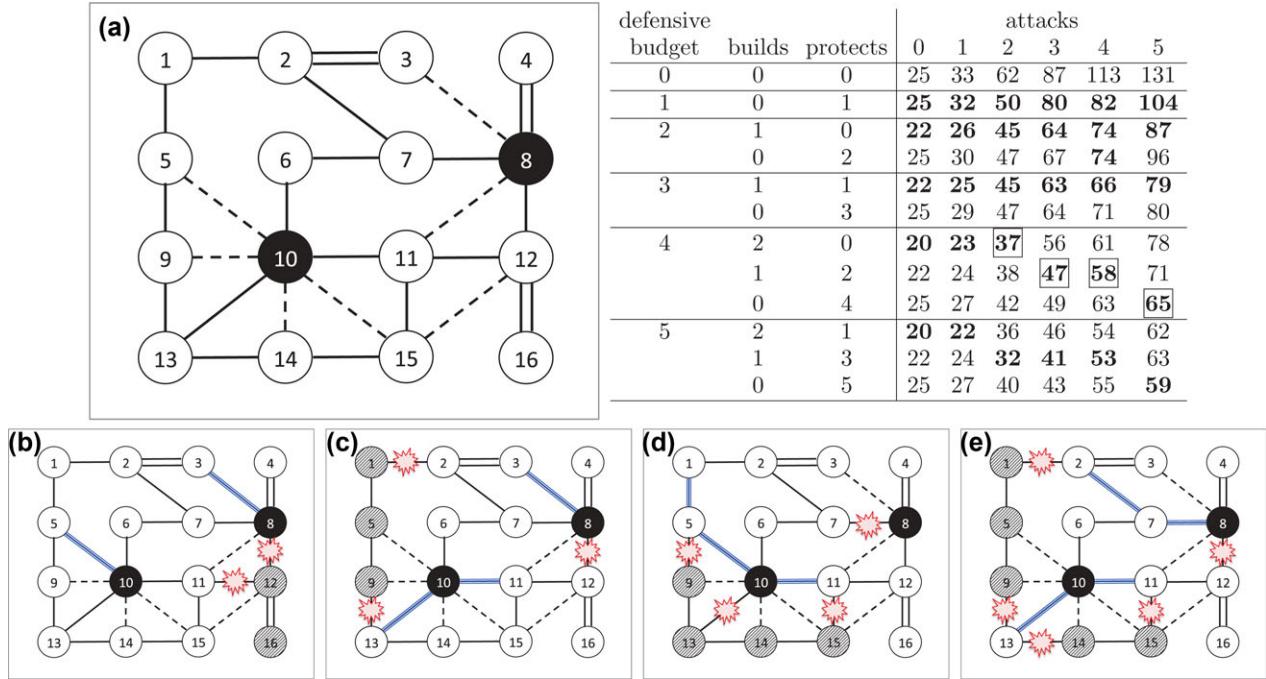


Fig. 7. Improving operational resilience with new construction. Panel (a) shows the potential (dashed) links that are available for construction. Under the assumption that building a new link costs twice as much as protecting an existing one, we consider the optimal defensive investment for different budget levels. The table shows the postattack operating cost for 0–5 attacks; values in bold correspond to the optimal (lowest-cost) defensive investments for the given budget and the specified number of attacks. The values in boxes correspond to the cases in Panels (b)–(e), showing the optimal investment of a defensive budget of four when there are 2–5 attacks, respectively. In some situations, it is better to build new links, while in others it is better to protect existing ones. In this example, the defenses *cannot* be prioritized into a rank-ordered list.

defense solutions when the defense budget is four units. An optimal defense against two attacks (Fig. 7(b)) is to build two new links, specifically [3, 8] and [5, 10]. An optimal defense against three attacks (Fig. 7(c)) is to build one new link ([3,8]) and then protect two links ([10,11] and [10,13]). An optimal defense against four attacks (Fig. 7(d)) is to build a different new link ([5,10]) and then protect a different pair of links ([1,5] and [10, 11]). An optimal defense against five attacks (Fig. 7(e)) is to protect four links ([2, 7], [7, 8], [10, 11], and [10, 13]). Thus, the best combinations of links to build or protect can be very different depending on the number of attacks.

4.3.3. Committing to a Defense

Casting infrastructure resilience in terms of our Attacker Model (Equation (2)) and Defender Model (Equation (3)) allows us to identify the sets of component losses that result in worst-case operating costs, as well as the defenses (via protection or new construction) that optimally mitigate these worst-case disruptions. However, as shown with our

notional infrastructure system, what is “best” in terms of defense often depends specifically on the number of attacks, and thus the links in our example cannot be strictly prioritized into a simple rank-ordered list. This is ubiquitous because the value of a component depends on its interaction with others (see Alderson *et al.*⁽⁷⁾ for a discussion).

In general, we will not know the size of the disruption that we will face. The point of this is that by presenting the resilience of the system in terms of a curve, one does not make any judgments *a priori* about the specific disruption magnitudes that are relevant. Uncertainty about the actual magnitude of disruption that we face is mitigated by showing the sensitivity of the system to different levels of disruption magnitude.

Nonetheless, decisions about defensive investment, particularly when they involve physical construction that is permanent, require that we commit to a single defensive plan, often articulated as a priority list and perhaps implemented in stages over time. Given the necessity to select only a single defense, we can solve for a prioritized list of components to

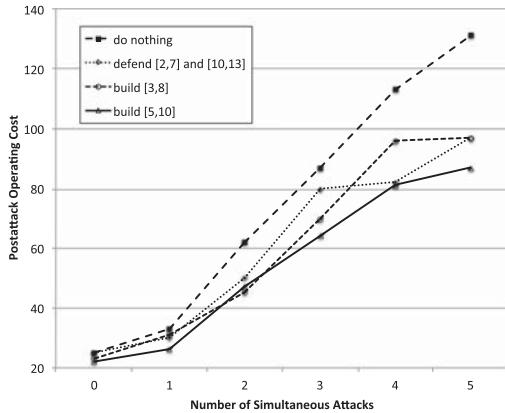


Fig. 8. Example of nondominant investment options. The resilience curves that would result from three defensive investments—namely (1) protect links [2, 7] and [10, 13], (2) build link [3, 8], and (3) build link [5, 10]—show that none of these is strictly “more resilient” than the others. However, given these choices, building link [5, 10] seems to be the best choice.

protect in an iterative manner. Specifically, we enumerate the number of defenses (i.e., $defense_tbudget = 1, 2, \dots$), and solve for a single new defense at each step and then fix the defense variable corresponding to that defended component for subsequent steps.

More simply, in the context of our notional infrastructure, assume that we have a defensive budget of two units, meaning for this system that we can either build a single new link or that we can protect two links. Based on the results in Fig. 4, we observe that links [2, 7] and [10, 13] are among the most important to protect. Also, our analysis of new construction reveals that links [3, 8] and [5, 10] are among the most important to build. Given these three defense options, which one is the best, and how does it compare to the status quo?

Fig. 8 illustrates the resilience curve for the baseline system, along with the resilience curves that would result from each of these three possible defensive investments. We observe that each of these three options results in postattack operating costs that are strictly lower than the baseline system, meaning that any of these defensive investments would yield a system that is more resilient than the current one. However, we also observe that none of the new resilience curves is strictly lower than the others, meaning that none of these solutions dominates the others in terms of the resilience that it provides. Nonetheless, we observe that building link [5, 10] yields the lowest postattack operating cost for all cases except for two

attacks, and even there it is a close second choice. For this reason, building link [5, 10] seems to be the best defensive investment decision for this assumed budget based on resilience as the only criteria.

In practice, real defensive investment decisions are likely to depend on not just resilience and cost, as described here, but also other regulatory, economic, and political criteria. Further, the example shows that the resilience curves for different investment options might not strictly dominate one another, making it impossible to say that one system is more resilient than another.⁽⁷⁾ Even so, the use of resilience curves to quantify operational resilience is a critically important first step toward more rigorous cost-benefit analysis for infrastructure defense.

5. DISCUSSION OF MODELS

While it might not be possible to specify *a priori* the final consequence associated with any disruption, the ability to articulate the Operator Model as a set of rules, operating procedures, or as a normative decision model means that it is possible to explore “what-if” scenarios using numerical simulation or other techniques.

In the models presented here, we restrict attention to a system of components, and in doing so we narrow the view of possible disruptions to those that are known to affect the function of the infrastructure. Moreover, when assessing the operational resilience of an infrastructure system, we restrict the adaptive behavior of the system to the choices defined in the Operator Model. Some researchers have argued that a key feature of resilience is the ability of a system to reconfigure itself in the presence of disruption (i.e., to self-organize, as might be done by humans in response to a disaster). To the extent that one can describe the way in which this might happen, it becomes possible to incorporate this in the Operator Model. If one is unable to describe either this emergent behavior or the rules that might lead to it, then our approach to resilience suffers no more than any other in predicting system response to disruption.

Nonetheless, the models in this article have been deliberately restricted in scope to keep our analysis simple and accessible. We therefore comment on ways in which the techniques presented in this article can be adapted to consider a broader range of issues.

5.1. Model Scalability

The notional infrastructure in this article is deliberately small for illustrative purposes, but the type

of Operator Model presented here can be solved at very large scale. In the commercial world, companies routinely solve problems involving millions of variables and hundreds of thousands of constraints. The bilevel Attacker Model is typically at least an order of magnitude or more larger than the Operator Model to which it is applied because of the need to solve for the best flows in response to every combination of attacks under consideration. Even here, these models can be solved at large scale. For example, when considering worst-case disruptions to an electric power grid, Salmerón *et al.*^(23,84) have solved Attacker Models consisting of thousands of electrical buses, high-voltage lines, transformers, and substations. Defender Models are more complicated yet because of the interaction of defense and attack combinations. Nonetheless, using appropriate decomposition algorithms,^(83,85) the trilevel optimization in the Defender Model can be solved very efficiently—all of the computational results in this article take less than 30 minutes to generate on a laptop computer.

We have formulated and solved Defender Models significantly larger than the example presented in this article. One such realistic model that we have built and solved considers the traffic handling function of the major roads and bridges of the San Francisco Bay Area.⁽⁸⁶⁾ The Operator Model contains 91 nodes and 266 directed arcs, represents origin-destination traffic demands between every pair of 2,292 census tracts, includes an extremely accurate piece-wise-linear approximation to the nonlinear congestion function currently used by California Department of Transportation traffic engineers,⁽⁸⁷⁾ and has been validated against rush-hour traffic patterns under many actual scenarios, including the loss of the San Francisco Bay Bridge during major repairs. Results from this model include (1) the discovery that the blockage of a small section of Interstate 880 for a single day can cause more short-term disruption to commuting traffic than the complete closure of any of the seven major bridges over the same time, and (2) the loss of the Bay Bridge for two years is more disruptive than the loss of the Golden Gate Bridge for five years. Insights from this model have been cited directly by the DHS,⁽⁸⁸⁾ and the Defender Model, which explicitly models increased operational costs (such as delays and reduced capacity) for protecting bridges, tunnels, and highway segments, solves within a few hours and provides clear defensive plans that are both face-valid, mathematically sound, and politically defensible.

5.2. Model Extensions

The constrained optimization problems described here are general mathematical programs, and as such, the example in this article can be extended in any number of ways. For example, we can replace the simple constraints on the number of attacks and the number of defenses with more general constraints on attacker and defender capabilities, including resource constraints that express, for example, limits on finances, personnel, or equipment, logical constraints on permitted or prohibited combinations of targets. That is, if we have a reasonable estimate of the resources that constrain adversary behavior (money, labor, explosives, etc.), and we have reasonable estimates of the cost of each attack in terms of each of these resources, then we can write one budget constraint for each resource, and end up with a model that considers those details, but is no harder to solve than the original. Examples of such constraints have been applied to large capital planning models⁽⁸⁹⁾ and models of industrial projects.⁽⁹⁰⁾ We can have as many of these constraints as we desire; they do not complicate the models significantly, and our solution algorithms remain effective. With more general constraints on attacker capability we can, for instance, model defense options that make attacks more expensive. (To a limited extent, we have already done this. Attacking the parallel links has a cost of two, and we assume that new construction costs twice that of protection. Again, see the appendices for details.)

Thus, we can also model *deterrence*—that is, the “stay at home” behavior of an attacker whose costs have been rendered intolerable by our defenses.

In addition, our models can easily incorporate *persistence* in defenses. That is, if, for any reason, we are committed to protecting a subset of components, we can fix the associated defensive variables in our Defender Model (Equation (3)). And just as for attacks, we can identify not just the optimal course action for the defender, but also the second-best (or third-best, etc.) defenses and their relative benefit. Such an enumeration of near-optimal courses action allows a policymaker to reconcile tradeoffs between the quality of a defense and other factors not explicitly represented in the model (e.g., political or social acceptability).

Here, we have restricted attention to the complete loss (or defense) of components. This technique generalizes to accommodate the partial loss (or defense), but we do not describe that here. We can

also model other “shocks” to the system, for example, a dramatic change in the demands placed on the system, as happens for regional transportation when there is a mass evacuation.

Although building an Operator Model requires significant up-front investment, the relatively slow pace at which infrastructure changes means that these models are available for reuse when the need arises.

The model in this article considers only a single instantaneous attack and a single instantaneous response to that attack. We can also model the behavior of the system over time, including the repair (or reconstitution) of components that return to operation on some forecast schedule. However, we typically do not model multiple attacks over time because we assume that after the first attack, the operating conditions change substantively enough for both the operator and the attacker to preclude other attacks over our planning horizon. Here, we are modeling singular events, not full-scale war.

5.3. Model Applicability

The modeling and analysis techniques described in this article have been applied to a variety of systems. Most relevant to this article, these techniques have been applied to real infrastructure systems across a range of sectors. In the context of the electric power grid, attacker–defender models have been used to analyze the vulnerability of major portions of the U.S. national grid,⁽⁹¹⁾ the dependence of U.S. military installations on the public grid,⁽⁹²⁾ and the resilience of electric infrastructure in U.S. territories.⁽⁹³⁾ The basic technique has also been successfully applied to civilian and military petroleum pipeline systems^(4,5,94,95) and multimodal transport of petroleum and coal.^(96–98) The use of constrained optimization to improve operations or system restoration has also been applied to natural gas infrastructure systems.^(99,100) Attacker–defender techniques have been applied to telecommunication systems, specifically terrestrial backbone networks,⁽¹⁰¹⁾ undersea cable systems,⁽¹⁰²⁾ and wireless networks.⁽¹⁰³⁾ Defender–attacker–defender techniques have been successfully applied to regional highway transportation systems^(83,86) and railroad systems.^(7,104)

The use of constrained optimization and game theory for identifying worst-case disruptions to operations and for planning defenses against them has applicability to more than just infrastructure systems.

Again, the key to the successful application of this technique is the development of operational models of system behavior. Recent success stories include the study of worst-case adversarial action in the context of industrial projects,⁽⁹⁰⁾ undersea warfare,⁽¹⁰⁵⁾ and ballistic missile defense.⁽¹⁰⁶⁾ These are examples where the Operator Model does not take the form of a network flow problem. We have built dozens of operational models of various infrastructures, each with their own peculiarities, and so far we have not found any that cannot be modeled in some reasonable way. If the Operator Model can be formulated and solved in a reasonable amount of time then the formulations of the Attacker Model and the Defender Model are usually straightforward, and the algorithms to solve them are now standard.⁽⁸⁵⁾

If the Operator Model is nonconvex, or if it is nonlinear and contains discrete variables, the models might take significantly longer to solve, or might require a linear or quadratic approximation. However, the formulations for the Attacker and Defender Models would still follow the same pattern. It is even possible to use simulation-optimization, where the Operator Model is itself a simulation, and the Attacker and Defender Models use optimization wrapped around this. The algorithms to solve these models have to be adapted a bit, and might end up being more heuristic, but the technique is general.⁽⁸⁵⁾

5.4. The Role of Uncertainty

The mathematical formulations here are deterministic, in the sense that all model inputs are assumed with certainty, and the “result” of any single model excursion follows directly from those inputs. In practice, we plan on solving many model excursions with different inputs. This type of parametric analysis can be of much greater practical value than the classical sensitivity analysis taught in optimization textbooks (see Brown and Rosenthal⁽¹⁰⁷⁾ for a discussion).

Although our focus in this article is on worst-case disruptions to infrastructure operation, our Operator Model is agnostic about the source of a disruption. In the realm of natural disasters, accidents, or random failures, we might try to define a probability distribution over the set of disruptive events X , and replace the worst-case (“max”) operator from the Attacker Problem with, for example, an expectation or some other measure of risk.

In practice, the expected value is often a poor choice of measure for risk-informed decisions.^(108,109)

Our point here is simply that one can use probabilistic techniques for risk in conjunction with our Operator Models, for those who favor these techniques. For simplicity, we restrict attention to calculating and minimizing the expected disruption.

The appropriate form of the expected disruption, as formulated here, is a stochastic optimization problem:

$$\mathbb{E}_{\tilde{x}} \left[\min_{\mathbf{y} \in Y(\tilde{x})} f(\tilde{x}, \mathbf{y}) \right]. \quad (4)$$

Here, \tilde{x} is a random variable and $\mathbb{E}_{\tilde{x}}$ denotes the expectation with regard to \tilde{x} . As with Equation (2), the operator takes action only after the (now random) disruption, and thus Equation (4) represents the expected cost of operating the system in the presence of disruption. The validity of such calculations hinges entirely on estimates of the probability distribution for \tilde{x} . We are wary of such estimates, particularly when they involve correlations between system components, and we therefore choose to focus exclusively here on the admittedly (and deliberately) conservative max-min formulation.

Defending against random disruptions becomes no more complicated. We seek defensive investment \mathbf{w} to minimize the expected cost of operating the system in the presence of a disruption:

$$\min_{\mathbf{w} \in W} \mathbb{E}_{\tilde{x}} \left[\min_{\mathbf{y} \in Y(\mathbf{w}, \tilde{x})} f(\mathbf{w}, \tilde{x}, \mathbf{y}) \right]. \quad (5)$$

As long as the Operator Model is formulated as an optimization problem (see Birge and Louveaux⁽¹¹⁰⁾ for an introduction to formulating and solving stochastic optimization models), our models and algorithms can be applied with no significant change.

In practice, infrastructure system owners and operators must contend with both expected and worst-case disruptions, and in principle a combination of Equations (3) and (5) could be used to obtain the required insight. Such ideas have been considered by Zhuang and Bier⁽¹¹¹⁾ in the context of “intentional and unintentional threats,” but their innermost models are not sufficiently “operational” to study infrastructure in the way we have described here.

5.5. Interdiction Versus Hijacking

The technique in this article works well for situations involving interdiction of system components. But it implicitly assumes that components are either going to be present and functional, or absent.

This assumption is sometimes known as “fail off” in the context of communication systems, and it has been an underlying assumption for the architectural design of the Internet.^(112,113) However, a very different situation arises when the system has components that “fail on”—that is, they continue to interact with other system components, but do not follow the rules, or *protocols*, for interaction. This type of disruption can lead to system *hijacking*, that is, the system continues to operate but behaves in a way that is not intended.^(17,113) Instances of hijacking are prevalent in technological and biological systems, and they represent some of the most challenging problems in these domains because it is sometimes the very mechanisms designed to create robustness and resilience that are hijacked for other purposes.^(14,17,114)

The techniques in this article are not designed to assess the impact of hijacking. Nonetheless, the types of disruptions considered here account for a large number of possible scenarios, and addressing them would go a long way to making infrastructure systems more resilient. Handling these types of hijacking scenarios, particularly as they pertain to cyber vulnerabilities, is an important topic for future research.

5.6. Robust Optimization

There is now a growing literature in the field of *robust optimization* that dates back to Wald’s min-max model for worst-case uncertainty.^(115,116) Robust optimization has been applied to a variety of problems in discrete optimization and network flows.^(117,118) Most of these models take a bilevel form—there is an initial design stage followed by the realization of an uncertain scenario. In the context of our infrastructure defense problems, this corresponds to a *defender–attacker* problem,⁽⁵⁾ in which the defender makes an initial investment in hardening or prepositioning, and the attacker follows with the worst-case attack. Our trilevel Defender Model can be viewed as a type of robust optimization in which there is an additional inner model of operation (after the uncertain attack is realized) that includes adaptation. This inclusion of adaptation through the use of an Operator Model is what distinguishes our models from the existing literature in robust optimization.

6. CONCLUSION

To introduce and demonstrate our definition of system resilience, along with supporting analytic

techniques, we intentionally chose a model instance that is so simple the reader can grasp normal operations by inspection. Yet, for the same example it is not easy to answer straightforward questions about how the system operator would respond to damage, how interdependencies between components yield vulnerabilities that can seriously disrupt system function, or how the defender should allocate limited resources to increase resilience to damage. This is where having a validated mathematical model of system operation offers tremendous value—it can provide a rapid and objective calculation of the consequence of damage to any set of components, and can therefore be used to identify vulnerabilities and to evaluate the improvement in resilience provided by any defensive plan.

The United States is currently spending billions of dollars on homeland security via federal, state, and local governments, and the most recent policy guidance in PPD21 suggests that resilience is going to be a key objective in future spending. Given this large investment, we strongly advocate the use of methods that (1) reflect the operation of an infrastructure as a system and evaluate its continuity of function in the presence of a disruptive event, (2) incorporate the inherent ability of existing infrastructure systems to adapt to disruptions or changes to their operating environment, and (3) facilitate the systematic exploration of disruptive events and their potential consequences, whether or not they are perceived as likely threats.

Our definition of resilience is qualitatively consistent with suggestions that have been made in the past, including by our most senior government policymakers, but we also show how to make quantitative assessments and evaluate specific alternatives for real systems. These techniques scale up to realistic size and fidelity,^(91–100) and admit a host of standard models, many already in use by system operators. We have used scores of these models to

assess resilience of a wide range of systems. Again and again, the same insights emerge: (1) the ability to assess actual system function is the key to an objective evaluation of consequence, (2) systems consist of individual components, but these components interact in complex ways and usually cannot be evaluated in isolation, (3) simple rank ordering of actions by any player is usually impossible, (4) trying to guess what an attacker might do instead of systematically evaluating his feasible courses of action underestimates vulnerability, and overestimates resilience, and (5) it is important to have a definition of resilience that is unambiguous and relies on well-documented, reproducible modeling and computation. We have been able to present our resilience assessments to senior policymakers at the local, state, and federal levels, with confidence that they fully understand our analysis, and we have frequently seen this advice implemented to good effect.

ACKNOWLEDGMENTS

This research was supported by the Office of Naval Research, the Air Force Office of Scientific Research, and the Defense Threat Reduction Agency. The authors gratefully acknowledge Kevin Wood and Javier Salmerón for ongoing discussions and collaborations that have contributed to the ideas in this article.

APPENDIX A: OPERATOR MODEL MATHEMATICAL FORMULATION

Although the example in this article is simple enough that the base flows can be solved by inspection, we present the formal Operator, Attacker, and Defender Models needed to obtain complete results.

In what follows, we use *barrels* as fuel units and *dollars* as cost units, but this is generic.

Indices and Sets

| | |
|---|---|
| $n \in N$ | nodes (alias i, j) |
| $[i, j] \in E$ | undirected edge between nodes i and j , $i < j$ |
| $(i, j) \in A$ | directed arc from node i to node j |
| $[i, j] \in E \iff (i < j) \wedge ((i, j) \in A \wedge (j, i) \in A)$ | |

Data [units]

| | |
|----------------|--|
| c_{ij} | per unit cost of traversing arc $(i, j) \in A$ [dollars/barrel] |
| u_{ij} | upper bound on total (undirected) flow on edge $[i, j] \in E$ [barrels] |
| \hat{x}_{ij} | 1 if edge $[i, j] \in E$ damaged, 0 otherwise [binary] |
| q_{ij} | per unit penalty cost of traversing arc $(i, j) \in A$ if damaged [dollars/barrel] |
| d_n | fuel supply at node $n \in N$ [barrels] |
| | (-demand for $d_n < 0$) |
| p_n | per unit penalty cost for demand shortfall $n \in N$ [dollars/barrel] |

Decision Variables [units]

| | |
|----------|--|
| Y_{ij} | flow on arc $(i, j) \in A$ [barrels] |
| S_n | fuel shortfall at node $n \in N$ [barrels] |

Formulation

$$\min_{Y, S} \sum_{[i, j] \in E} [(c_{ij} + q_{ij}\hat{x}_{ij}) Y_{ij} + (c_{ji} + q_{ji}\hat{x}_{ij}) Y_{ji}] + \sum_{n \in N} p_n S_n \quad (D0)$$

$$\text{s.t. } \sum_{(n, j) \in A} Y_{nj} - \sum_{(i, n) \in A} Y_{in} - S_n \leq d_n \quad \forall n \in N \quad (D1)$$

$$0 \leq Y_{ij} + Y_{ji} \leq u_{ij} \quad \forall [i, j] \in E \quad (D2)$$

$$S_n \geq 0 \quad \forall n \in N \quad (D3)$$

Discussion

The objective function (D0) combines the total flow cost and the total penalty cost. Constraints (D1) enforce balance of flow at each node. Stipulations (D2) and (D3) ensure bounds on decision variables. This formulation implements *cost-based interdiction*—that is, damage to an arc makes it extremely expensive but not infeasible—which makes the problem easier to solve computationally.

In the above example, we have $d_n = 10$ for $n \in \{8, 10\}$ and $d_n = -1$ otherwise. In addition, we set $c_{ij} = 1$, $u_{ij} = 15$, and $q_{ij} > 10$ for all $(i, j) \in A$, $u_{ij} \geq 14$ for all $[i, j] \in E$, and $p_n = 10$ for all $n \in N$.

APPENDIX B: ATTACKER MODEL MATHEMATICAL FORMULATION

The Attacker Model builds on the previous formulation but has additional elements.

Additional Data [units]

r_{ij} “cost” to break edge $[i, j] \in E$ [cardinality]

$attack_budget$ budget constraint on the number of simultaneous attacks [cardinality]

Additional Decision Variables [units]

X_{ij} 1 if attacker breaks edge $[i, j] \in E$, 0 otherwise [binary]

Formulation

$$\begin{aligned} \max_{X, Y, S} \quad & \sum_{[i, j] \in E} [(c_{ij} + q_{ij}X_{ij})Y_{ij} + (c_{ji} + q_{ji}X_{ij})Y_{ji}] + \sum_{n \in N} p_n S_n \\ \text{s.t.} \quad & \text{(D1), (D2), (D3)} \end{aligned} \quad (\text{AD0})$$

$$\sum_{[i, j] \in E} r_{ij}X_{ij} \leq attack_budget \quad (\text{AD1})$$

$$X_{ij} \in \{0, 1\} \quad \forall [i, j] \in E \quad (\text{AD2})$$

The objective function (AD0) is the same as that for the Operator Model (D0), except that parameters \hat{x}_{ij} have been replaced by decision variables X_{ij} . Constraint (AD1) limits the number of simultaneous attacks, and the cost to attack each edge can be different. Stipulations (AD2) require that attacks are binary. We note that $q_{ij} = 0$ implies that arc (i, j) is effectively *invulnerable* because attacking it does not increase the flow cost for the operator.

In the above example, we model parallel edges as costing twice as much to attack. That is, we have $r_{2,3} = r_{4,8} = r_{12,16} = 2$ and all other $r_{ij} = 1$.

APPENDIX C: DEFENDER MODEL MATHEMATICAL FORMULATION

The Defender Model builds on the previous formulation but has additional elements.

Additional Sets

| | |
|------------------|---|
| E^B | set of additional edges available to be built, $E^B \cap E = \emptyset$ |
| $[i, j] \in E^B$ | $\iff (i < j) \wedge ((i, j) \in A \wedge (j, i) \in A)$ |

Additional Data [units]

| | |
|-------------------|---|
| h_{ij} | “cost” to protect edge $[i, j] \in E$ [cardinality] |
| h_{ij}^B | “cost” to build edge $[i, j] \in E^B$ [cardinality] |
| $defense_budget$ | budget constraint on the number of defenses [cardinality] |

Additional Decision Variables [units]

| | |
|------------|---|
| W_{ij} | 1 if defender protects edge $[i, j] \in E$, 0 otherwise [binary] |
| W_{ij}^B | 1 if defender builds edge $[i, j] \in E^B$, 0 otherwise [binary] |

Formulation

$$\begin{aligned}
 \min_{W, W^B} \max_{X} \min_{Y, S} \quad & \sum_{[i, j] \in E} [(c_{ij} + q_{ij} X_{ij}(1 - W_{ij})) Y_{ij} + (c_{ji} + q_{ji} X_{ij}(1 - W_{ij})) Y_{ji}] + \\
 & \sum_{[i, j] \in E^B} [c_{ij} Y_{ij} + c_{ji} Y_{ji}] + \sum_{n \in N} p_n S_n \\
 \text{s.t.} \quad & (D1), (D2), (D3), (AD1), (AD2) \\
 & 0 \leq Y_{ij} + Y_{ji} \leq u_{ij} W_{ij}^B \quad \forall [i, j] \in E^B \quad (DAD1) \\
 & \sum_{[i, j] \in E} h_{ij} W_{ij} + \sum_{[i, j] \in E^B} h_{ij}^B W_{ij}^B \leq defense_budget \quad (DAD2) \\
 & W_{ij} \in \{0, 1\} \quad \forall [i, j] \in E \quad (DAD3) \\
 & W_{ij}^B \in \{0, 1\} \quad \forall [i, j] \in E^B \quad (DAD4)
 \end{aligned} \tag{DAD0}$$

The objective (DAD0) includes the cost of flow over existing (and possibly damaged or protected) edges, flow over newly built edges, and penalties for unmet demand. Constraints (DAD1) allow flow only on new edges if they have been built. The constraint (DAD2) requires that the cost of all defenses fall within the existing defense budget; the cost of protecting and building edges can be different. Stipulations (DAD3) and (DAD4) enforce binary defenses.

In the above example, we assume that it costs twice as much to build a new edge as to protect an existing one, that is, $h_{i,j} = 1, \forall [i, j] \in E$ and $h_{i,j}^B = 2, \forall [i, j] \in E^B$.

In principle, solving the Defender Model requires nothing more than enumerating every possible combination of defense and attack, then solving the corresponding Operator Model for each, and then finding the one that yields the lowest cost. In practice, such enumeration is impractical. Alderson *et al.* (83,85) provide details of a decomposition algorithm to solve models of this type without exhaustive enumeration.

REFERENCES

1. The White House. Presidential Policy Directive: Critical Infrastructure Security and Resilience. Washington, DC, 2013.
2. Title 42 US Code, Sec 5195c et seq 2006 Supp IV, 2011. Critical Infrastructure Protection. Available at: <http://www.gpo.gov/>, Accessed February 14, 2015.
3. Homeland Security Council (HSC). National Strategy for Homeland Security. Washington, DC: White House, 2007.
4. Brown G, Carlyle W, Salmerón J, Wood R. Analyzing the vulnerability of critical infrastructure to attack, and planning defenses. Pp. 102–123 in Greenberg H, Smith J (eds). Tutorials in Operations Research: Emerging Theory, Methods, and Applications. Hanover, MD: Institute for Operations Research and Management Science, 2005.
5. Brown G, Carlyle WM, Salmerón J, Wood K. Defending critical infrastructure. Interfaces, 2006; 36:530–544.
6. Wood AJ, Wollenberg BF. Power Generation, Operation and Control, 2nd ed. New York: Wiley, 1996.
7. Alderson D, Brown G, Carlyle W, Cox L. Sometimes there is no “most vital” arc: Assessing and improving the operational resilience of systems. Military Operations Research, 2013; 8:21–37.
8. Shapley L. A value for n -person games. Pp. 307–317 in Kuhn H, Tucker A (eds). Contributions to the Theory of Games, Vol. II, Volume 28 of Annals of Mathematical Studies. Princeton, NJ: Princeton University Press, 1953.

9. Ford L, Fulkerson D. Maximal Flow Through a Network. Technical Report. Santa Monica, CA: RAND Corporation, Research Memorandum RM-1400, 1954.
10. Alderson D. Catching the “network science” bug: Insight and opportunity for the operations researcher. *Operations Research*, 2008; 56:1047–1065.
11. Albert R, Albert I, GL N. Structural vulnerability of the North American power grid. *Physical Review E*, 2004; 69:025103–025106.
12. Wang JW, Rong LL. Cascade-based attack vulnerability on the US power grid. *Safety Science*, 2009; 47:1332–1336.
13. Hines P, Cotilla-Sanchez E, Blumsack S. Do topological models provide good information about electricity infrastructure vulnerability? *Chaos: An Interdisciplinary Journal of Nonlinear Science*, 2010; 20:033122–033122.
14. Doyle JC, Alderson D, Li L, Low S, Roughan M, Shalunov S, Tanaka R, Willinger W. The “robust yet fragile” nature of the Internet. *Proceedings of the National Academy of Sciences of the United States of America*, 2005; 102:14497–14502.
15. Willinger W, Alderson DL, Doyle JC. Mathematics and the Internet: A source of enormous confusion and great potential. *Notices of the AMS*, 2009; 56:586–599.
16. Department of Homeland Security (DHS). Critical Infrastructure Sector Partnerships, 2013. Available at: <http://www.dhs.gov/critical-infrastructure-sector-partnerships>, Accessed August 3, 2013.
17. Alderson D, Doyle J. Contrasting views of complexity and their implications for network-centric infrastructures. *IEEE Transactions on Systems, Man, Cybernetics A: Systems and Humans*, 2010; 40:839–852.
18. O’Neill R, Helman U, Hobbs B, Baldick R. Independent system operators in the United States: History, lessons learned, and prospects. Pp. 479–528 in Sioshansi, F, Pfaffenberger, W (eds). *Electricity Market Reform: An International Perspective*. Oxford: Elsevier, 2006.
19. Wardrop JG. Some theoretical aspects of road traffic research. *Proceedings of the Institute of Civil Engineers, Part II*, 1952; 1:325–378.
20. Beckmann M. On the theory of traffic flows in networks. *Traffic Quarterly*, 1967; 2:109–116.
21. Rardin R. *Optimization in Operations Research*. Upper Saddle River, NJ: Prentice Hall, 1997.
22. Alderson D, Li L, Willinger W, Doyle J. Understanding Internet topology: Principles, models, and validation. *IEEE/ACM Transactions on Networking*, 2005; 13:1205–1218.
23. Salmerón J, Wood K, Baldick R. Analysis of electric grid security under terrorist threat. *IEEE Transactions on Power Systems*, 2004; 19:905–912.
24. Hwang CL, F A Tillman F, Lee MH. System-reliability evaluation techniques for complex/large systems: A review. *IEEE Transactions on Reliability*, 1981; R-30:416–423.
25. Billinton R, Allan R. *Reliability Evaluation of Engineering Systems: Concepts and Techniques*, 2nd ed. New York: Plenum Press, 1992.
26. Jorion P. *Value at Risk: A New Benchmark for Measuring Derivatives Risk*. New York: Irwin Professional Publishers, 1996.
27. Rockafellar RT, Uryasev S. Conditional value-at-risk for general loss distributions. *Journal of Banking & Finance*, 2002; 26:1443–1471.
28. Page M, Alderson D, Doyle J. The magnitude distribution of earthquakes near southern California faults. *Journal of Geophysical Research: Solid Earth*, 2011; 116:1978–2012.
29. Hiemer S, Jackson DD, Wang Q, Kagan YY, Woessner J, Zechar J, Wiemer S. A stochastic forecast of California earthquakes based on fault slip and smoothed seismicity. *Bulletin of the Seismological Society of America*, 2013; 103:799–810.
30. Paté-Cornell M, Guikema S. Probabilistic modeling of terrorist threats: A systems analysis approach to setting priorities among countermeasures. *Military Operations Research*, 2002; 7:5–23.
31. Garrick B, Hall J, McDonald JC, OToole T, Probst PS, Parker E, Rosenthal R, Trivelpiece A, Van Arsdale L, Zebroski E. Confronting the risks of terrorism: Making the right decisions. *Reliability Engineering and System Safety*, 2004; 86:129–176.
32. Parnell G, Liebe R, Dillon-Merrill R, Buedo D, Scouras J, Colletti B, Cummings M, McGarvey D, Newport R, Vinch P. *Homeland Security Risk Assessment: Volume I—An Illustrative Framework and Volume II: Appendices of Methods*. Washington, DC: Homeland Security Institute, 2005.
33. Willis HH, Morral AR, Kelly TK, Medby JJ. *Estimating Terrorism Risk*. Santa Monica, CA: Rand Corporation, 2006.
34. McGill W, Ayyub B, Kaminskiy M. Risk analysis for critical asset protection. *Risk Analysis*, 2007; 27:1265–1281.
35. Ezell BC, Bennett SP, Von Winterfeldt D, Sokolowski J, Collins AJ. Probabilistic risk analysis and terrorism risk. *Risk Analysis*, 2010; 30:575–589.
36. Willis HH. Guiding resource allocations based on terrorism risk. *Risk Analysis*, 2007; 27:597–606.
37. ASME (American Society of Mechanical Engineers) Innovative Technologies Institute. RAMCAP, Risk Analysis and Management for Critical Asset Protection, 2008. Available at: <http://www.asme-iti.org/RAMCAP>, Accessed May 14, 2011.
38. Keeney R. Modeling values for anti-terrorism analysis. *Risk Analysis*, 2007; 27:585–596.
39. Bier V. Choosing what to protect. *Risk Analysis*, 2007; 27:607–620.
40. Bier VM, Haphuriwat N, Menoyo J, Zimmerman R, Culpen AM. Optimal resource allocation for defense of targets based on differing measures of attractiveness. *Risk Analysis*, 2008; 28:763–770.
41. Department of Homeland Security (DHS). National Infrastructure Protection Plan. Washington, DC: Department of Homeland Security, 2009.
42. National Research Council (NRC). Committee on Methodological Improvements to the Department of Homeland Security’s Biological Agent Risk Analysis. *Department of Homeland Security Bioterrorist Risk Assessment: A Call for Change*. Washington, DC: National Academies Press, 2008.
43. National Research Council (NRC). Committee to Review the Department of Homeland Security’s Approach to Risk Analysis. *Review of the Department of Homeland Security’s Approach to Risk Analysis*. Washington, DC: National Academies Press, 2010.
44. Cox A. Some limitations of risk = threat × vulnerability × consequence for risk analysis of terrorist attacks. *Risk Analysis*, 2008; 28:1749–1761.
45. Cox A. What’s wrong with hazard-ranking systems? An expository note. *Risk Analysis*, 2009; 29:940–948.
46. Cox L. Game theory and risk analysis. *Risk Analysis*, 2009; 29:1062–1068.
47. Brown G, Cox A. How probabilistic risk assessment can mislead terrorism risk analysis. *Risk Analysis*, 2011; 31:196–204.
48. Brown G, Cox A. Making terrorism risk analysis less harmful and more useful: Another try. *Risk Analysis*, 2011; 31:193–195.
49. Apostolakis GE, Lemon DM. A screening methodology for the identification and ranking of infrastructure vulnerabilities due to terrorism. *Risk Analysis*, 2005; 25:361–376.
50. Stamatelatos M, Dezfuli H, Apostolakis G, Everline C, Guarro S, Mathias D, Mosleh A, Paulos T, Riha D, Smith C, Vesely W, Youngblood R. *Probabilistic Risk Assessment Procedures Guide for NASA Managers and Practitioners*, 2011. Technical Report NASA/SP-2011-3421, NASA.

51. Wood R. Bilevel network interdiction models: Formulations and solutions. Pp. 1-11 in Cochran J (ed.) Wiley Encyclopedia of Operations Research and Management Science. New York: John Wiley & Sons, 2011. [doi10.1002/9780470400531.eorms0932]
52. Washburn A, Wood R. Two-person zero-sum games for network interdiction. *Operations Research*, 1995; 43:243–251.
53. von Stackelberg HV. Grundlagen einer reinen Kostentheorie. Vienna: Verlag von Julius Springer, 1932.
54. Park J, Seager T, Rao P, Convertino M, Linkov I. Integrating risk and resilience approaches to catastrophe management in engineering systems. *Risk Analysis*, 2013; 23:356–367.
55. Holling C. Resilience and stability of ecological systems. *Annual Review of Ecology and Systematics*, 1973; 4:1–23.
56. Hollnagel E, Woods D, Leveson N (eds). Resilience Engineering: Concepts and Precepts. Aldershot, UK: Ashgate Press, 2006.
57. Madni A, Jackson S. Towards a conceptual framework for resilience engineering. *IEEE Systems Journal*, 2009; 3:181–191.
58. Haimes YY. On the definition of resilience in systems. *Risk Analysis*, 2009; 29:498–501.
59. Westrum R. A typology of resilience situations. Pp. 49–60 in Hollnagel E, Woods D, Leveson N (eds). Resilience Engineering: Concepts and Precepts. Aldershot, UK: Ashgate Press, 2006.
60. Zolli A, Healy A. Resilience: Why Things Bounce Back. New York: Free Press, 2012.
61. Hale A, Heijer T. Defining resilience. Pp. 95–123 in Hollnagel E, Woods D, Leveson N (eds). Resilience Engineering: Concepts and Precepts. Aldershot, UK: Ashgate Press, 2006.
62. Woods D. Essential characteristics of resilience. Pp. 49–60 in Hollnagel E, Woods D, Leveson N (eds). Resilience Engineering: Concepts and Precepts. Aldershot, UK: Ashgate Press, 2006.
63. Leveson N, Dulac N, Zipkin D, Cutcher-Gershenfeld J, Carroll J, Barrett B. Engineering resilience into safety-critical systems. Pp. 95–123 in Hollnagel E, Woods D, Leveson N (eds). Resilience Engineering: Concepts and Precepts. Aldershot, UK: Ashgate Press, 2006.
64. Ottino J. Engineering complex systems. *Nature*, 2004; 427:399.
65. Willinger W, Govindan R, Jamin S, Paxson V, Shenker S. Scaling phenomena in the Internet: Critically examining criticality. *Proceedings of the National Academy of Sciences USA*, 2002; 99:2573–2580.
66. Haimes Y, Crowther K, Horowitz B. Homeland security preparedness: Balancing protection with resilience in emergent systems. *Systems Engineering*, 2006; 11:287–308.
67. Reed D, Kapur K, Christie R. Methodology for assessing the resilience of networked infrastructure. *IEEE Systems Journal*, 2009; 3:174–180.
68. Ta C, Goodchild A, Pitera K. Structuring a definition of resilience for the freight transportation system. *Transportation Research Record: Journal of the Transportation Research Board*, 2009; 2097:19–25.
69. Chen L, Miller-Hooks E. Resilience: An indicator of recovery capability in intermodal freight transport. *Transportation Science*, 2012; 46:109–123.
70. Nair R, Avetisyan H, Miller-Hooks E. Resilience framework for ports and other intermodal components. *Transportation Research Record: Journal of the Transportation Research Board*, 2010; 2166:54–65.
71. Omer M, Mostashari A, Nilchiani R, Mansouri M. A framework for assessing resiliency of maritime transportation systems. *Maritime Policy & Management*, 2012; 39:685–703.
72. Freckleton D, Heaslip K, Louisell W, Collura J. Evaluation of resiliency of transportation networks after disasters. *Transportation Research Record: Journal of the Transportation Research Board*, 2012; 2284:109–116.
73. Hughes J, Healy K. Measuring the resilience of transport infrastructure. Technical Report Research Report 546. Washington, DC: Transportation Research Board, the National Academies, 2014.
74. Vugrin E, Warren D, Ehlen M, Camphouse R. A framework for assessing the resilience of infrastructure and economic systems. Pp. 77–116 in Gopalakrishnan K, Peeta S (eds). Sustainable and Resilient Critical Infrastructure Systems: Simulation, Modeling, and Intelligent Engineering. New York: Springer-Verglag, 2010.
75. Vugrin E, Warren D, Ehlen M. A resilience assessment framework for infrastructure and economic systems: Quantitative and qualitative resilience analysis of petrochemical supply chains to a hurricane. Pp. 77–116 in Gopalakrishnan D, Peeta S (eds). Proceedings of 6th Global Congress on Process Safety, American Institute of Chemical Engineers. San Antonio, TX, 2010.
76. Rose A. Defining and measuring economic resilience to disasters. *Disaster Prevention and Management*, 2004; 13:307–314.
77. Rose A. Economic resilience to natural and man-made disasters: Multidisciplinary origins and contextual dimensions. *Environmental Hazards*, 2007; 7:383–398.
78. Rose A. Economic Resilience to Disasters. Technical Report, CARRI Research Report 8, Oakridge, TN: Community & Regional Resilience Institute, 2009.
79. US Government Accountability Office (GAO). Critical Infrastructure Protection: Update to National Infrastructure Protection Plan Includes Increased Emphasis on Risk Management and Resilience, 2010, Washington, DC: US Government Accountability Office Report, GAO-10-296, March 5, 2010.
80. Department of Homeland Security (DHS), Risk Steering Committee. DHS Risk Lexicon. Washington, DC, 2010.
81. Flynn S. The Edge of Disaster: Rebuilding a Resilient Nation. New York: Random House, 2007.
82. Brown G, Dell R. Formulating linear and integer linear programs: A rogues' gallery. *INFORMS Transactions on Education*, 2007; 7:153–159.
83. Alderson D, Brown G, Carlyle W, Wood RK. Solving defender-attacker-defender models for infrastructure defense. Pp. 28–49 in Wood K, Dell R (eds). Operations Research, Computing and Homeland Defense. Hanover, MD: Institute for Operations Research and the Management Sciences, 2011.
84. Salmerón J, Wood K, Baldick R. Worst-case interdiction analysis of large-scale electric power grids. *IEEE Transactions on Power Systems*, 2009; 24:96–104.
85. Alderson D, Brown G, Carlyle W. Assessing and Improving Operational Resilience of Critical Infrastructures and Other Systems. Hanover, MD: Institute for Operations Research and Management Science, 2014 [to appear in *Tutorials in Operations Research*].
86. Alderson D, Brown G, Carlyle W, Wood RK. Optimizing the Operational Resilience of Regional Infrastructure: A Case Study of the Highway System in the San Francisco Bay Area, Presentation, INFORMS Computing Society Meeting, Santa Fe, NM, January 6, 2013.
87. California Metropolitan Transportation Commission. Initial Examination of Volume Delay Functions Using PeMS Data, 2012. Available at: http://mtcgis.mtc.ca.gov/foswiki/pub/Main/Documents/2012_03_06_RELEASE_Volume_delay_functions.pdf, Accessed April 30, 2013.
88. Department of Homeland Security (DHS), Homeland Infrastructure Threat and Risk Analysis Center. Infrastructure Impact Assessment 28 October 2009: San Francisco Bay Bridge Closure. Washington, DC: DHS, 2009.
89. Brown G, Dell R, Newman A. Optimizing military capital planning. *Interfaces*, 2004; 34:415–425.

90. Brown G, Carlyle W, Harney R, Skroch E, Wood R. Interdicting a nuclear-weapons project. *Operations Research*, 2009; 57:866–877.
91. Salmerón J, Wood K. Final Report on DOE Research Project DE-AI02-05ER25670: Reducing the Vulnerability of Electric Power Grids to Terrorist Attack. Technical Report NPS-OR-09-003-PR, Naval Postgraduate School. Distribution authorized to U.S. Government Agencies only (sensitive information), 2009.
92. Salmerón J, Alderson D, Brown G. Resilience Report: Electric Power Infrastructure Supporting Mission Assurance at Vandenberg Air Force Base (U). Naval Postgraduate School, Technical Report NPS-OR-11-008, 2011. [Distribution authorized to U.S. Government Agencies and their contractors due to military infrastructure.]
93. Salmerón J, Alderson D, Brown G, Wood R. Resilience Report: The Guam Power Authority Electric Power Grid: Analyzing Vulnerability to Physical Attack (U). Naval Postgraduate School, Technical Report NPS-OR-12-002, 2012. [Distribution authorized to DoD and DoD contractors only due to infrastructure vulnerability analysis.]
94. Chankij MK, Assessing the Resiliency of the JP8 Distribution System on Guam. Master's thesis, Naval Postgraduate School, Monterey, CA, 2012.
95. Montgomery JD. Oahu Petroleum Infrastructure Resilience. Master's thesis, Naval Postgraduate School, Monterey, CA, 2013.
96. Onuska J. Defending the Pittsburgh Waterways Against Catastrophic Disruption. Master's thesis, Naval Postgraduate School, Monterey, CA, 2012.
97. Burton C. Analyzing the U.S. Military Fuel Distribution Network on Okinawa. Master's thesis, Naval Postgraduate School, Monterey, CA, 2013.
98. Long C. Analyzing the Resilience of the Fuel Distribution System for Mainland Japan. Master's thesis, Naval Postgraduate School, Monterey, CA, 2013.
99. Avery W, Brown G, Rosenkranz J, Wood R. Optimization of purchase, storage and transmission contracts for natural gas utilities. *Operations Research*, 1992; 40:446–462.
100. Coffrin C, van Hentenryck P, Bent R. Last mile restoration for multiple interdependent infrastructures. In *Proceedings of the Twenty-Sixth AAAI Conference on Artificial Intelligence (AAAI 2012)*, Toronto, Canada, 2012.
101. Barkley T. An Attacker-Defender Model for IP-Based Networks. Master's thesis, Naval Postgraduate School, Monterey, CA, 2007.
102. Crain J. Assessing Resilience in the Global Undersea Cable Infrastructure. Master's thesis, Naval Postgraduate School, Monterey, CA, 2012.
103. Shankar A. Optimal Jammer Placement to Interdict Wireless Network Services. Master's thesis, Naval Postgraduate School, Monterey, CA, 2008.
104. Babick JP. Tri-Level Optimization of Critical Infrastructure Resilience. Master's thesis, Naval Postgraduate School, Monterey, CA, 2009.
105. Brown G, Kline J, Thomas A, Washburn A, Wood K. A game-theoretic model for defense of an oceanic basin against submarines. *Operations Research*, 2011; 59:25–40.
106. Brown G, Carlyle M, Diehl D, Kline J, Wood K. A two-sided optimization for theater ballistic missile defense. *Operations Research*, 2005; 53:263–275.
107. Brown G, Rosenthal R. Optimization tradecraft: Hard-won insights from real-world decision support. *Interfaces*, 2008; 38:356–366.
108. Aven T, Kørte J. On the use of risk and decision analysis to support decision-making. *Reliability Engineering & System Safety*, 2003; 79:289–299.
109. Savage S. *The Flaw of Averages: Why We Underestimate Risk in the Face of Uncertainty*. New York: John Wiley & Sons, 2009.
110. Birge JR, Louveaux F. *Introduction to Stochastic Programming*. New York: Springer, 1997.
111. Zhuang J, Bier V. Balancing terrorism and natural disasters: Defensive strategy with endogenous attack effort. *Operations Research*, 2007; 55:976–991.
112. Clark DD. The design philosophy of the DARPA Internet protocols. *ACM SIGCOMM Computer Communications Review*, 1988; 18:106–114.
113. Doyle JC, Carlson J, Low SH, Paganini F, Vinnicombe G, Willinger W, Parrilo P. Robustness and the Internet: Theoretical foundations. In Jen E (ed). *Robust Design: A Repertoire from Biology, Ecology, and Engineering*. Oxford, UK: Oxford University Press, 2003.
114. Doyle J, Csete M. Architecture, constraints, and behavior. *Proceedings of the National Academy of Sciences*, 2011; 108:15624–15630.
115. Wald A. Statistical decision functions which minimize the maximum risk. *Annals of Mathematics*, 1945; 46:265–280.
116. Danskin J. *The Theory of Max-Min*. New York: Springer-Verlag, 1967.
117. Bertsimas D, Sim M. The price of robustness. *Operations Research*, 2004; 52:35–53.
118. Bertsimas D, Brown D, Caramanis C. Theory and applications of robust optimization. *SIAM Review*, 2011; 53:464–501.